Universal Quantum Tracking and Infinite Compute: What the Full Wavefunction History Would Let Us Infer About Laws, Causality, and Emergence

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 16, 2026

Imagine a thought experiment in which every particle in the universe carries an inert identifier and, more radically, the complete quantum state of the universe is available for analysis at all times: phases, amplitudes, and entanglement structure—an effective record of $t \mapsto |\Psi(t)\rangle$. Pair this omniscient dataset with infinite compute. What could be inferred, not by philosophical preference, but by algorithmic reconstruction from the total history itself? This paper treats "physics" as an inverse problem: the laws of nature are sought as the minimal generator capable of reproducing the observed universal quantum trajectory. We ask what becomes decidable—unitary versus non-unitary dynamics, the form of the Hamiltonian or quantum channel, locality and causal structure, conservation laws and symmetries, and the emergence of classical records and thermodynamic arrows from entanglement growth. We then examine what remains underdetermined even with perfect data: competing ontologies related by dualities or subsystem factorizations, the impossibility of proving absolute minimality of a compression, and the persistent gap between description and ultimate "why." The goal is a disciplined map of inference limits and powers when measurement is no longer the bottleneck—only interpretation and computability are.

Keywords: universal wavefunction, quantum state tomography, infinite compute, Laplace's demon, Hamiltonian reconstruction, unitary evolution, collapse models, Lindblad dynamics, entanglement structure, decoherence, causal inference, spacetime emergence, quantum gravity, renormalization group, effective field theory, thermodynamic arrow of time, algorithmic information theory, Kolmogorov complexity, model underdetermination, subsystem factorization.

A collaboration with GPT-5.2-Thinking. 97 pages. CC4.0.

Outline

1. Introduction: From Laplace's Demon to Universal Quantum Tracking
1.1 The classical demon and the quantum obstacle
1.2 Upgrading the thought experiment: full access to $t \mapsto |\Psi(t)\rangle$
1.3 What counts as "deduction" versus "interpretation"
1.4 The paper's central claim: laws as the best compression of history

2. Formalizing the Dataset: What "Complete Quantum Access" Means
2.1 The object of study: pure state history $|\Psi(t)\rangle$ and mixed histories $\rho(t)$
2.2 Factorizations and partitions: defining subsystems without smuggling in ontology
2.3 Observables as derived, not primitive: correlators, marginals, and records
2.4 Practical idealizations: continuity in time, resolution, and error-free access

3. Reconstructing Dynamics: Inferring the Generator of Time Evolution
3.1 From snapshots to propagators: reconstructing $U(t2,t1)$ or quantum channels
3.2 Continuous-time inference: extracting $H(t)$ or the most general generator
3.3 Constraints from composition, locality, and symmetry
3.4 Identifiability: when multiple generators reproduce the same history

4. Deciding Unitarity: Collapse, Noise, or Pure Decoherence
4.1 Global purity tests and reversibility as empirical criteria
4.2 Distinguishing environmental decoherence from objective collapse
4.3 Parameter inference for collapse models: rates, length scales, and couplings
4.4 What "Born rule" becomes when the universal state is fully known

5. Symmetries and Conservation Laws as Algorithmic Invariants
5.1 Noether structure recovered from the inferred generator
5.2 Gauge redundancy versus physical symmetry: what the data can and cannot fix
5.3 Detecting symmetry breaking and phase transitions from $|\Psi(t)\rangle$
5.4 Hidden sectors: inferring conserved quantities tied to unknown degrees of freedom

6. Locality and Causality: Building the Universe's Influence Graph
6.1 Causal discovery in quantum systems: interventions without interventions
6.2 Light cones, finite propagation, and effective causal speed limits
6.3 Separating entanglement correlations from causal influence
6.4 Diagnosing nonlocal dynamics: what would count as a genuine violation

7. Spacetime and Geometry: Fundamental Stage or Emergent Description
7.1 Reconstructing geometry from trajectories versus from entanglement structure
7.2 When a metric is the best compression: information geometry criteria
7.3 Curvature, horizons, and global topology in the full-state record
7.4 What would falsify “spacetime is emergent” in this setting

8. Emergence of the Classical World: Records, Robustness, and Pointer Structures
8.1 What a “classical fact” is: stable records in subsystems
8.2 Decoherence as entanglement bookkeeping: exact mutual-information flows
8.3 Redundant encoding and objectivity: why many observers agree
8.4 The boundary of classicality: where classical description stops compressing well

9. The Arrow of Time: Exact Accounting of Entropy and Irreversibility
9.1 Entanglement entropy versus thermodynamic entropy: aligning the ledgers
9.2 Low-entropy boundary conditions and the growth of accessible correlations
9.3 Coarse-graining choice as a physical question: which macrostates are natural
9.4 Could an arrow emerge without special initial conditions? What the data would show

10. The Renormalization Ladder: Deriving Effective Laws at Every Scale
10.1 Integrating out degrees of freedom as a controlled compression
10.2 Exact EFT emergence: why “effective” becomes precise in this experiment
10.3 Phase transitions, universality classes, and fixed points as features of $|\Psi(t)\rangle$
10.4 Bridging to chemistry and biology: what counts as a legitimate derivation

11. Counterfactuals and Contingency: What “Would Have Happened” Can Mean
11.1 From one realized history to a lawful ensemble: the counterfactual leap
11.2 Sensitivity analysis: micro-perturbations and macro-outcomes
11.3 Control knobs and bottlenecks: rare degrees of freedom with outsized leverage
11.4 Limits: counterfactuals under dualities and non-unique factorizations

12. Underdetermination and Duality: When Perfect Data Doesn’t Pick One Story
12.1 Dual theories, identical predictions: the compression-equivalence problem
12.2 Subsystem factorization ambiguity: “what is a system?” as a choice
12.3 Ontology versus invariants: what remains fixed across equivalent descriptions
12.4 A practical rule: prefer invariants and operational reconstructions

13. Computability Boundaries: Even Infinite Compute Doesn’t Prove Minimality
13.1 Laws as shortest programs: the appeal and the trap
13.2 Why absolute Kolmogorov minimality is not generally decidable
13.3 What can be proven instead: dominance, stability, and model class minimality
13.4 A disciplined epistemology: compression as guide, not as oracle

14. The Residual “Why”: Where Explanation Stops Being Deduction
14.1 Why these laws rather than others: boundary conditions and meta-laws
14.2 Fine-tuning, selection, and anthropic structure: what the dataset can constrain
14.3 Meaning, value, and teleology: inference limits and legitimate questions
14.4 A modest conclusion: maximal description still leaves metaphysics intact

15. Conclusion: The Inference Frontier with Complete Quantum History
15.1 What becomes rigidly inferable: dynamics, symmetries, causality, emergence
15.2 What remains plural: ontology, factorization, ultimate grounds
15.3 The main takeaway: the universe as data, laws as compression, humility as method
15.4 Future work: formal pipelines for generator discovery and emergence auditing

16. References

Art

1. Introduction: From Laplace's Demon to Universal Quantum Tracking

The dream of total explanation has always worn two faces. One face is predictive: if we knew everything about the state of the world at a given moment and possessed enough computational power, could we compute the entire future? The other face is reconstructive: if we could see the world's total history at sufficient resolution, could we infer what the world is and what governs it? Laplace's famous demon is usually read as the first face—an omniscient calculator turning the present into the future. The thought experiment developed here shifts the emphasis toward the second: the universe as a complete dataset, and the laws of nature as what an ideal analyst would recover from that dataset.

We begin with the classical demon because it expresses a tempting clarity: world + state + law = trajectory. But quantum theory complicates that equation at the level where "knowing the state" and "reading the state" are not conceptually separable. The obstacle is not mere technological difficulty; it is structural. In ordinary quantum mechanics, perfect tracking is forbidden by the fact that measurement disturbs, that noncommuting observables cannot be jointly specified, and that entanglement makes subsystem states partial and contextual. If we insist on staying within ordinary operational constraints, Laplace's demon becomes a fable about impossible access.

This paper therefore makes an explicit "upgrade." We assume a physically inert tagging mechanism and a godlike readout that provides, for all times, the full quantum state of the universe. With that move, the limiting factor is no longer measurement but inference: what can be deduced from a total state history, and what remains interpretive even when the data are complete?

1.1 The classical demon and the quantum obstacle

Laplace's demon begins with a premise that seems innocuous in classical mechanics: at an instant, the world has definite values for all relevant degrees of freedom, and those values can be known without changing them. In that framework, the demon needs two ingredients: (i) the complete classical microstate (positions and momenta for all particles, or field values and conjugate momenta), and (ii) the equations of motion. Under these assumptions, the future and past are, in principle, fixed by the present. The universe is a deterministic dynamical system, and the demon is merely a powerful integrator.

Quantum mechanics disrupts each piece of that premise.

First, the notion of "the complete state" is different. The quantum state is not a list of simultaneously well-defined values of all observables. It is a vector (or density operator) that

assigns probabilities and phases to potential measurement outcomes. The state encodes amplitude structure, not classical facts. If one insists on reading it as “hidden classical data,” one is already stepping into interpretation.

Second, the act of “knowing” the state is structurally constrained. The quantum obstacle is not simply that we lack better microscopes; it is that measurement is not a passive look. Noncommuting observables cannot be simultaneously measured with arbitrary precision, and generic measurement changes the system. Even if one imagines measuring everything, the very attempt would scramble the thing one wanted to preserve. The classical demon presupposes a clean separation between ontology (what is) and epistemology (what is known). Quantum theory entangles the two.

Third, entanglement undermines the local bookkeeping the classical demon relies on. In classical physics, the state of a part is a restriction of the state of the whole. In quantum physics, the reduced state of a subsystem is generally mixed and does not capture all correlations with the rest of the universe. The world resists being decomposed into independently well-defined parts. A demon that only monitors “particles” as if they were separable classical objects will miss what is most characteristically quantum: relational phase information and nonlocal correlation structure.

Finally, there is the measurement problem. Standard quantum theory gives unitary evolution for isolated systems and nonunitary “collapse” (or update) upon measurement. But what counts as a measurement if the demon is “measuring everything”? The classical demon assumes one universal rule. In quantum theory, the presence of two rules (unitary evolution and collapse) creates ambiguity about whether the demon’s omniscience is even definable without committing to an interpretation.

So the quantum obstacle is not merely “hardness.” It is a formal barrier: in ordinary quantum mechanics, “tagging and tracking without disturbance” is not obviously coherent for generic quantum degrees of freedom, because access is coupled to dynamics. If we want a demon-like scenario, we must either accept a particular interpretation that supplies an underlying determinate microhistory, or we must modify the thought experiment to grant access to what ordinary observers cannot access.

That is what we do next.

1.2 Upgrading the thought experiment: full access to $t \mapsto |\Psi(t)\rangle$

The upgraded scenario is deliberately maximalist. We assume that every fundamental degree of freedom is perfectly identified, and—more strongly—that we can read out the full universal quantum state at all times. Symbolically, we are given the entire trajectory:

$t \mapsto |\Psi(t)\rangle$

where $|\Psi(t)\rangle$ denotes the universal pure state. If one prefers a more general setting (e.g., objective noise, collapse, or effective openness), we can allow a density operator history $t \mapsto \rho(t)$. But the core idea is the same: the dataset is the universe's complete quantum description through time, including all phase information and all entanglement structure.

This upgrade changes the problem in three decisive ways.

First, it dissolves the operational measurement barrier by stipulation. We do not ask how such access is physically implemented, because the point of the thought experiment is not engineering; it is inference. The usual objection—"you can't measure the wavefunction without disturbing it"—is set aside. The upgraded demon is not an experimenter; it is an analyst with perfect access to the ontic state history (whatever that turns out to mean).

Second, it makes the universe's history the primary object, not isolated experiments. In ordinary science, we infer laws from a sparse sample of situations: we run experiments, vary inputs, and model outcomes. Here, the dataset is not sparse. It is total. The analytic task becomes closer to system identification from complete logs than to hypothesis testing from a small number of trials. You do not ask "what happens if we do X?" first; you ask "what generator produces all of this?"

Third, it shifts the central difficulty from prediction to explanation. If you are literally handed $|\Psi(t)\rangle$ for all t—past and future—then "predicting what will happen" is trivial in the sense that the future is already in the record. The deeper question becomes: what compact structure makes this immense object intelligible? What features of the history are lawful regularities and what are contingent details? What is the simplest dynamical description that reproduces the entire trajectory?

This is the key methodological pivot of the paper: we treat fundamental laws as the most powerful compression of the universe's complete history. The demon is no longer merely a calculator that forecasts; it is an algorithmic physicist that reconstructs dynamics, symmetries, causal structure, and emergent layers of description from an all-encompassing state history.

1.3 What counts as “deduction” versus “interpretation”

When the dataset is maximal, it becomes easy to speak too quickly. One can say, “With complete $|\Psi(t)\rangle$, everything is deducible.” But that conflates at least three different senses of “deduce,” and the differences matter.

First is **logical entailment** from the data alone. This is the strongest sense: a conclusion is deduced if no alternative is possible given the dataset. In this sense, many metaphysical claims are not deducible even from perfect data, because different ontological stories can generate exactly the same state history, or can redescribe it without changing any invariant content.

Second is **model identification within a chosen class**. If we restrict attention to a family of candidate dynamical laws (say, local Hamiltonians on a fixed Hilbert space factorization, or Markovian quantum channels of a certain form), then complete state history can uniquely identify parameters and sometimes the form of the generator. But that uniqueness depends on the chosen model class. It is “deduction under assumptions,” not deduction in the absolute sense.

Third is **compression-based inference**. Here we choose the “best” theory as the one that minimizes description length: the shortest program that reproduces $|\Psi(t)\rangle$ to the required precision. This approach has immense explanatory appeal because it operationalizes Occam’s razor. But it is not the same as strict entailment. Multiple theories may tie for best compression, and proving global minimality runs into deep computability limits. Compression is a disciplined heuristic for theory choice, but it does not magically turn metaphysics into arithmetic.

So we need a clean line. In this paper:

- “Deduction” will mean recoverable structure that is stable under reasonable re-descriptions and that follows from the dataset within clearly specified inferential rules: the generator of time evolution (up to equivalence), the presence or absence of nonunitarity, the symmetry and conservation structure encoded in invariants, and the causal influence patterns implied by the dynamics.
- “Interpretation” will mean choices of ontology, factorization, and explanatory language that are not uniquely forced by $|\Psi(t)\rangle$: whether one regards $|\Psi\rangle$ as the whole of reality or as a statistical object over deeper variables; whether one treats branching as “many worlds” or as a calculational tool; whether spacetime is fundamental or emergent if both descriptions fit equally well; and, most importantly, whether the recovered regularities imply meaning, purpose, or teleology.

This distinction is not a dismissal of interpretation. It is a commitment to intellectual hygiene. If the thought experiment is meant to clarify what maximal knowledge buys, then we must also clarify what maximal knowledge does not automatically buy.

1.4 The paper's central claim: laws as the best compression of history

The central claim can be stated simply: given complete access to the universal quantum state history, the fundamental laws of physics are the minimal dynamical generator that reproduces that history, and the hierarchy of "sciences" above physics are successively coarser compressions that preserve stable predictive structure.

This is both a methodological and a philosophical claim.

Methodologically, it reframes discovery. Instead of guessing laws and testing them against sparse data, we solve an inverse problem: find the operator dynamics (unitary or nonunitary) that maps $|\Psi(t1)\rangle$ to $|\Psi(t2)\rangle$ for all times, and choose among candidate representations by favoring those that compress the dataset most. Symmetries are recovered as invariances of the generator. Conservation laws are recovered as invariants. Locality is recovered as sparseness or constraint structure in the generator that limits influence propagation. Emergence is recovered as the existence of low-dimensional effective variables whose dynamics are stable under coarse-graining and that greatly reduce description length without sacrificing explanatory power.

Philosophically, it reframes "law" from a metaphysical primitive to an inferential optimum. Laws are not merely what we write down because they are elegant; they are what survive as the best explanation—meaning, the best compression—when the universe is treated as an enormous, information-rich object. In this view, to say "the universe obeys these laws" is to say "these laws are the most economical generator consistent with the entire history," where "economical" is not aesthetic but informational.

At the same time, the claim is bounded by humility. "Best compression" may not be unique. Dualities and gauge redundancies may permit multiple equally short descriptions. The minimal-program ideal has computability limits. And even if we recover a generator that reproduces everything, we have not thereby answered why that generator holds, why the initial condition is what it is, or why the whole exists at all.

But we have achieved something substantial: we have made precise what maximal physical knowledge would look like, what it could reconstruct, and where the boundary between physics and metaphysics remains even when measurement is no longer the bottleneck. That boundary—sharp, disciplined, and explicit—is the real payoff of the universal tracking thought experiment.

2. Formalizing the Dataset: What "Complete Quantum Access" Means

The strength of the thought experiment is also its danger: "complete quantum access" can be asserted in a sentence, but it hides a thicket of choices. What, exactly, is the data object? A pure state or a density operator? What is time, in a universe where time itself may be emergent? What does it mean to "tag every particle" in a setting where particles may be effective excitations rather than fundamental degrees of freedom? And perhaps most importantly: how do we speak about subsystems, observables, and records without quietly importing the very ontology we claim to infer?

This section therefore does something deliberately unglamorous: it puts the thought experiment on formal rails. We define the dataset in a way that is strong enough to support the inferences we want later (reconstructing dynamics, symmetries, causality, emergence), but cautious enough to avoid circularity. The guiding principle is this: we will treat the universal state history as primary, and treat most familiar objects—particles, fields, observers, measurements, classical facts—as derived structures extracted from that history via well-defined operations.

2.1 The object of study: pure state history $|\Psi(t)\rangle$ and mixed histories $\rho(t)$

In its most maximal form, the thought experiment hands us a function of time:

$t \mapsto |\Psi(t)\rangle$

where $|\Psi(t)\rangle$ is the state vector of the entire universe in a Hilbert space H. This is the "pure history" version. It implicitly assumes that the universe is closed and evolves in a way that preserves purity—typically, unitary evolution generated by some Hamiltonian (possibly time-dependent). In that setting, the dataset contains everything: not just probabilities for measurement outcomes, but the full amplitude and phase structure that supports interference and entanglement.

However, there are two reasons to broaden the formalism to include mixed histories.

First, the underlying dynamics might not be strictly unitary. If objective collapse, fundamental decoherence, or other nonunitary mechanisms exist, the universal state might be more naturally represented as a density operator $\rho(t)$. Even if one believes the universe is pure "underneath," the operationally available universal description in such theories may still be mixed. Since part of the paper's goal is to decide between unitary and nonunitary evolution using the dataset, it is methodologically cleaner to allow both representations at the outset.

Second, even within unitary quantum mechanics, the "state of a subsystem" is generally mixed because of entanglement with the rest of the universe. Later sections will rely heavily on reduced states, marginals, entropies, and mutual information. Those are naturally expressed in density-operator language. A clean formalism should make both pure and mixed descriptions first-class citizens.

So we define the dataset in the most general way as a time-indexed family of density operators:

$t \mapsto \rho(t)$

with the understanding that the "pure history" case corresponds to $\rho(t) = |\Psi(t)\rangle\langle\Psi(t)|$ at all times.

This still leaves an important question: what is the time parameter t? In ordinary physics, t is a coordinate (or proper time along a foliation). In quantum gravity or emergent-time scenarios, "time" may not be fundamental. For this paper's purpose, we adopt a pragmatic stance: the dataset is indexed by whatever parameter plays the role of dynamical ordering for the universal evolution. If later the best compression of the dataset suggests that time is emergent from correlations, then "t" will be recognized as a convenient label for slices of a deeper structure rather than a metaphysical primitive. But we do not need to decide that here. We only require that the dataset supports consistent composition of evolution: moving from t1 to t2 to t3 should agree with moving directly from t1 to t3, in whatever generalized sense applies.

Finally, "complete access" means not merely access to expectation values or measurement statistics, but direct access to the full operator $\rho(t)$ itself. In ordinary quantum mechanics that is not operationally available without ensembles and disturbance. Here it is given by stipulation. This is exactly what makes later inference possible: if you know $\rho(t)$ completely, you can compute any derived quantity without further measurement.

2.2 Factorizations and partitions: defining subsystems without smuggling in ontology

The phrase "tag every particle" suggests a classical picture: the universe is made of identifiable objects, and each object can be labeled and tracked. But quantum theory, especially in field-theoretic and emergent descriptions, warns us that "particles" are not always fundamental. They can be excitations of fields; they can be frame-dependent; they can be approximate and context-sensitive. If we import "particle ontology" into the formalization, we risk baking in what we later want to infer.

The clean way to proceed is to treat the universe as a single quantum system with Hilbert space H, and treat subsystems as decompositions or partitions of that structure. In the simplest textbook case, one writes:

$H = H_A \otimes H_B$

for two subsystems A and B. More generally, we may have many factors:

$H = \bigotimes_i H_i$

and partitions correspond to grouping subsets of factors.

But the key point is that factorizations are not unique, and not all physically meaningful decompositions look like simple tensor products (especially in quantum field theory, gauge theories, and gravity where local algebras and constraints complicate naive factorization). Even when tensor-product factorization exists, many different factorizations may be mathematically possible, and choosing one can implicitly choose an ontology: what counts as a "thing," where boundaries lie, which degrees of freedom are "local," and what counts as an "observer."

Therefore we adopt a careful stance:

- A "subsystem" is not assumed a priori as a set of particles.
- A subsystem is specified by a choice of factorization (or, more generally, by a choice of subalgebra of observables).
- Inferences about locality, causal structure, and emergent classicality must be robust across reasonable choices, or else explicitly conditional on the chosen decomposition.

This has two practical implications for the paper.

First, later reconstruction of "local dynamics" must be framed as: there exists a decomposition in which the generator is approximately local (e.g., sparse interactions among nearby degrees of freedom). If many inequivalent decompositions make the generator look nonlocal, then locality is not a straightforward property of the raw dataset; it is a property of a particular descriptive choice that may or may not be privileged by compression.

Second, "tagging" must be reinterpreted. The inert tag is best thought of as a way of fixing identity across time for degrees of freedom in a chosen decomposition, not as a metaphysical commitment to particle individuality. A tag is an index on factors (or modes) in the Hilbert space representation that stays consistent through the evolution. In classical language we would say "particle ID." In this formal language we should say "degree-of-freedom label relative to a chosen decomposition."

This reframing is not pedantry. If the project is to learn what can be inferred from $|\Psi(t)\rangle$, then we must be honest that some of what we call "the world's furniture" depends on how we choose to carve the Hilbert space. The dataset can constrain and privilege certain carvings (via compression and stability), but it does not automatically hand them to us for free.

2.3 Observables as derived, not primitive: correlators, marginals, and records

In everyday physics, observables are often treated as primitive: position, momentum, spin, field amplitude, and so on. But in this thought experiment, we have something stronger than "measurement outcomes." We have the state itself. That means observables should be treated as derived queries on the dataset rather than as fundamental givens.

Given ρ(t), any observable O (a Hermitian operator) has an expectation value:

$\langle O \rangle_t = \mathrm{Tr}[\rho(t)\, O]$

and any set of observables has correlators:

$\langle O_1 O_2 \ldots O_n \rangle_t = \mathrm{Tr}[\rho(t)\, O_1 O_2 \ldots O_n]$

The dataset therefore contains not just values of particular observables, but the capacity to answer every possible query of this form. In ordinary experimental science, access is limited: you can only measure a chosen subset, with disturbance and finite statistics. Here, correlators are computational outputs.

Reduced states are equally central. For a subsystem A, its reduced density operator is:

$\rho_A(t) = \mathrm{Tr}_{A^c}[\rho(t)]$

where the trace is over the complement of A. From these reduced states, we can compute entropies, mutual informations, and other information-theoretic quantities that diagnose entanglement structure and record formation:

$S(\rho_A) = -\mathrm{Tr}[\rho_A \log \rho_A]$
$I(A:B) = S(\rho_A) + S(\rho_B) - S(\rho_{AB})$

These derived quantities are not optional decorations; they are the main bridge from "full quantum access" to the familiar macroscopic world. In particular, the emergence of "records"—stable, redundant encodings that behave like classical facts—can be formalized as a pattern in these information measures. Roughly, a "record" exists when some subsystem B carries information about A in a way that is robust, redundantly distributed, and stable under typical perturbations. In ordinary language, many parts of the environment carry consistent information about a pointer basis of A, enabling multiple observers to agree. In our formal language, that agreement corresponds to a structured pattern of mutual information and conditional independence across partitions.

This is a crucial conceptual inversion: classical facts are not the input; they are the output. The data object is ρ(t). What we call "a measurement outcome" is a derived regularity: a stable pattern in the correlations between subsystems such that certain degrees of freedom serve as

"memory." A big part of the paper's later argument is that once you treat observables and records as derived, the measurement problem becomes less a mysterious rule and more a question of emergent structure: which decompositions and which bases generate stable records that compress well?

2.4 Practical idealizations: continuity in time, resolution, and error-free access

Even a thought experiment needs a discipline of idealizations. If we leave "complete access" undefined, any later inference can be criticized as relying on hidden assumptions. So we state the idealizations openly.

First, we assume **sufficient temporal resolution**. Inference of a generator (Hamiltonian or more general dynamical law) from a state history requires knowing how the state changes. If we only have widely spaced snapshots, many generators can interpolate them. In the limit of continuous access, one can infer instantaneous derivatives and thus reconstruct a local-in-time generator when such a generator exists. Therefore we idealize the dataset as continuous in t, or at least sampled densely enough that the relevant dynamical scales are resolved.

Second, we assume **sufficient state resolution**. "Knowing ρ(t)" is not merely knowing a few low-order moments; it is knowing the full operator to arbitrary precision. Many physical distinctions—especially those involving phase, interference, and fine entanglement structure—are invisible under coarse approximations. Since the thought experiment's purpose is to remove measurement limits, we assume this resolution is effectively infinite. If later we want to discuss approximate access, we can, but the core claims of maximal inferability assume ideal access.

Third, we assume **error-free access**. In practice, any data stream includes noise. But since the paper is exploring what is possible "in principle" with infinite compute and complete access, we take the data as exact. This allows us to discuss identifiability cleanly. If the dataset is exact and still does not uniquely fix some conclusion, then that conclusion is genuinely underdetermined, not merely hard to estimate.

Fourth, we assume **a stable identity map for degrees of freedom** relative to a chosen decomposition. "Tagging" is formalized as the ability to track which degrees of freedom at time t correspond to which degrees of freedom at time t+dt, so that the dataset is not a scrambled sequence of unrelated states. Without such consistency, the notion of dynamics is ill-posed. Again, we treat this as part of the stipulation: the dataset is coherently indexed.

Finally, we acknowledge a boundary: if the best description of the universe's history does not admit a global time parameter or a simple tensor factorization, then our formalism must be seen as a convenience rather than the final word. But the thought experiment's strength is that

it can reveal that. If the minimal compression of ρ(t) requires abandoning naive time or naive subsystem structure, the analysis will force that conclusion rather than hiding it.

In summary, “complete quantum access” is formalized as: a fully specified universal density-operator history ρ(t), indexed by an evolution parameter that supports composition, together with a coherent scheme for partitioning degrees of freedom into subsystems and computing any derived correlation, marginal, entropy, or information-flow quantity. Everything else—particles, measurements, classical facts, spacetime locality—is treated as an inferred structure that must earn its place by explanatory compression and stability.

3. Reconstructing Dynamics: Inferring the Generator of Time Evolution

If the dataset is the full universal state history, then the central inverse problem is unavoidable: what rule carries the state forward? In ordinary physics we start with a guessed form of law (Newton, Schrödinger, Einstein), then test it against limited observations. Here we invert the logic. We are given the trajectory and asked to infer the generator.

This move is conceptually clean but technically subtle. Even with perfect access to ρ(t), “the law” is not merely a curve-fitting exercise. We want a rule that is (i) consistent across all times, (ii) compositional (evolution from t1 to t3 equals evolution from t1 to t2 to t3), (iii) stable under small perturbations in how we represent subsystems, and (iv) compressive: it should reduce the informational burden of storing the entire history by replacing it with a compact generator plus initial condition (or plus boundary conditions).

This section shows how, in principle, one would reconstruct either a unitary propagator U(t2,t1) (if evolution is closed and unitary) or a general quantum channel (if evolution is nonunitary or effectively open). We then turn continuous-time reconstruction into a generator—Hamiltonian in the unitary case, or the most general dynamical generator (e.g., Lindblad-like) in the nonunitary case. Finally we address the main limitation: identifiability. Even perfect data do not guarantee a unique generator, because multiple generators can reproduce the same single realized trajectory. The question becomes: under what conditions does the dataset force uniqueness, and when does it merely constrain an equivalence class?

3.1 From snapshots to propagators: reconstructing U(t2,t1) or quantum channels

Start with the simplest formulation. Suppose we are handed exact states at discrete times:

ρ(t1), ρ(t2), ..., ρ(tN)

The most direct question is: what map takes ρ(t1) to ρ(t2)? In quantum theory, the most general physically admissible map on density operators is a completely positive, trace-preserving (CPTP) map—informally, a quantum channel. So we ask whether there exists a channel E(t2,t1) such that:

$\rho(t2) = E(t2,t1)[\rho(t1)]$

In the closed, unitary case, this channel has the special form:

$\rho(t2) = U(t2,t1)\ \rho(t1)\ U(t2,t1)^{\dagger}$

where U(t2,t1) is unitary.

Given a single pair (ρ(t1), ρ(t2)), there are infinitely many channels consistent with that input-output relation. The immediate reason is obvious: a channel is a function defined on all possible input states, but we only observe its action on one state. Even if we know the full universal state, we still only see one trajectory—one path through state space—not the full action of the map on a neighborhood.

So the reconstruction task must use the entire time series and, crucially, must exploit structural constraints: composition over multiple times, continuity, locality, and symmetry. The dataset is enormous in time, and that helps. Instead of one input-output pair, we have a long sequence of constraints:

$\rho(tk+1) = E(tk+1,tk)[\rho(tk)]$ for all k

If we assume time-homogeneity (or slowly varying structure), we can impose that these channels are instances of a smaller parametric family. Even without homogeneity, we can impose that the family composes consistently:

$E(t3,t1) = E(t3,t2) \circ E(t2,t1)$

This composition law is powerful because it couples many time steps. It prevents us from choosing an arbitrary channel at each step independently. In the unitary case, the analogous condition is:

$U(t3,t1) = U(t3,t2)\ U(t2,t1)$

and $U(t,t) = I$

With dense time sampling, these constraints become strong enough to support generator reconstruction in many cases, because small-step evolution becomes the primitive object.

There is another way to state the same idea that emphasizes what “complete access” buys. From ρ(t), we can compute not just the state, but any derived correlators for any chosen

partition. That means we can examine the evolution of reduced states for many subsystems simultaneously:

$\rho_A(t) = Tr_{A^c}[\rho(t)]$

and track how each $\rho_A(t)$ changes. If the universe admits a description in which subsystems interact locally, then the evolution of small subsystems will be constrained by a local generator acting on them and their boundaries. That gives many overlapping "views" of the global dynamics, dramatically increasing identifiability. In effect, we get a massive set of coupled input-output relations across scales and partitions, not just a single global relation.

So snapshot-to-propagator reconstruction is possible in principle, but only by leveraging the redundancy implicit in complete access: we are not limited to one observable; we can project the global state into countless subsystem perspectives and demand mutual consistency.

3.2 Continuous-time inference: extracting H(t) or the most general generator

Discrete propagators are not yet "the law." What we want is the generator of evolution: the object that, together with an initial state (or boundary conditions), compresses the entire history.

In the ideal limit, the dataset is continuous in time and exact. Then we can consider the time derivative of the state.

Unitary case (pure or mixed): Hamiltonian reconstruction

If evolution is unitary, a pure state evolves as:

$d|\Psi(t)\rangle/dt = -i\, H(t)\, |\Psi(t)\rangle$

and a density operator evolves as:

$d\rho/dt = -i\, [H(t), \rho(t)]$

Given $\rho(t)$ and $d\rho/dt$, this looks like a linear equation for $H(t)$. But there is an immediate caveat: for a single known trajectory $\rho(t)$, the commutator equation does not uniquely determine $H(t)$, because many different Hamiltonians can generate the same motion on that particular state. Any term that commutes with $\rho(t)$ at that time is invisible to the commutator. More precisely, if $K(t)$ satisfies $[K(t), \rho(t)] = 0$, then $H(t)$ and $H(t) + K(t)$ produce the same instantaneous derivative of $\rho(t)$.

So even in the unitary case, additional structure is needed to pin down H(t): constraints like locality (H is a sum of local terms), time smoothness, symmetry, and consistency across many reduced states.

Nonunitary case: master equation / quantum channel generator

If evolution is not unitary, the most general continuous Markovian form often used is:

$d\rho/dt = L_t(\rho)$

where L_t is a linear generator that preserves trace and complete positivity over finite times. In special cases, L_t can be represented in Lindblad form, with a Hamiltonian part plus dissipative terms. The paper does not need to assume Lindblad form universally; the point is broader: from continuous data ρ(t), one can infer whether the best-fitting generator is purely commutator (unitary) or requires additional terms.

The key diagnostic is structural: does there exist a representation in which the evolution of the global state is reversible and purity-preserving? If not, what minimal deviation (noise, collapse-like terms, nonlinearity) is required to reproduce the observed trajectory? Complete access lets us evaluate these questions without the usual measurement ambiguity because we are not inferring from sparse outcomes; we are examining the state itself.

In principle, continuous-time inference proceeds by fitting a generator that reproduces the observed derivative at each time while satisfying global constraints: compositionality over finite intervals, positivity, trace preservation, and whatever additional regularities (time smoothness, locality, symmetry) lead to maximal compression.

3.3 Constraints from composition, locality, and symmetry

The raw inverse problem is underconstrained if we only look at one trajectory. What makes “law discovery” possible is that plausible physical laws are not arbitrary maps: they have strong structural signatures. In the complete access setting, these signatures become computationally testable constraints.

Composition and consistency

The most basic constraint is that evolution composes. For any $t1 < t2 < t3$, the map from t1 to t3 must equal the composition of the maps $t1 \rightarrow t2$ and $t2 \rightarrow t3$. This rules out ad hoc step-by-step maps that do not fit into a coherent dynamical flow.

In continuous time, composition implies that the generator must be compatible with a well-defined time-ordered exponential (unitary case) or with a coherent family of channels (general

case). This consistency is not optional: without it, “dynamics” is merely a time-indexed lookup table, not a compressive law.

Locality as sparsity in the generator

If the world is local, the generator has constrained structure. In many frameworks, locality means the Hamiltonian decomposes as a sum of terms acting on small neighborhoods:

$H = \Sigma_x h_x$

or, in more general channel language, the generator L has a locality or finite-range property.

In the complete dataset, locality can be tested as a compression principle: does there exist a factorization/partition in which the inferred generator is sparse and local, and in which influence propagation is bounded? If yes, locality is not assumed; it is discovered as the representation that makes the generator simplest.

If no such representation exists—if every factorization that fits the data yields a generator with long-range, highly nonlocal couplings—then either the universe is nonlocal at a fundamental level or our chosen subsystem language is misaligned with the true degrees of freedom. In either case, “locality” becomes a derived and testable property, not an article of faith.

Symmetry as invariance of the generator

Symmetries are among the most compressive features of physical law because they impose repetition and constraint across space, time, and internal degrees of freedom. In the reconstructed dynamics, symmetry shows up as invariance of the generator under certain transformations. In unitary language, we look for transformations G such that:

H transforms into itself (or into an equivalent form) under G

and in channel language, the generator commutes with the symmetry action in the appropriate representation.

Because we have the entire state history, we can test symmetry not merely by checking whether certain outcomes repeat, but by checking whether the inferred generator must have symmetry to compress the data. If imposing a symmetry reduces description length dramatically without sacrificing exact reproduction of $\rho(t)$, then symmetry is not merely a human aesthetic preference; it is a property demanded by the structure of the history.

Composition, locality, and symmetry therefore function as the core triad of constraints that turns an underdetermined inverse problem into an inferable law: they restrict the hypothesis space to generators that are plausible as “laws” rather than arbitrary time-indexed explanations.

3.4 Identifiability: when multiple generators reproduce the same history

Even with perfect data, there is a hard conceptual fact: a single realized trajectory rarely identifies a unique dynamical law. The reason is not ignorance; it is mathematics. Many different generators can produce exactly the same path through state space, especially if you only demand agreement on that path rather than on all nearby counterfactual states.

There are several distinct sources of non-identifiability.

Gauge and representational freedom

Different mathematical descriptions can be physically equivalent. Gauge transformations, dualities, and reparameterizations can change the form of the generator without changing observable invariants. In the universal-state setting, this can appear as multiple generators related by a change of variables that leaves all derived predictions invariant. If two generators are related by such an equivalence, demanding a unique choice is a category mistake. The best we can do is identify the equivalence class and prefer members that are simplest under a chosen convention.

Commutant ambiguity on a single state

As noted above, in unitary evolution, any term that commutes with $\rho(t)$ is invisible in the instantaneous commutator equation. Over time, this can be partly broken (because what commutes at one time may not commute at another), but the general phenomenon remains: the state may fail to explore enough of Hilbert space to "excite" all directions that would reveal the full Hamiltonian.

Trajectory-only fit versus neighborhood fit

A generator is a rule acting on all states, but we only see its action on the states the universe actually visits. If the state visits a thin manifold in the full space of density operators, then the generator's behavior off that manifold can vary wildly without affecting the observed history. This is the deepest identifiability problem: perfect knowledge of one world does not automatically reveal the law of all possible worlds.

Breaking degeneracy via subsystem redundancy and minimality

Complete access gives us one powerful weapon against this: redundancy across partitions and scales. If we demand that a single generator simultaneously explains the evolution of all reduced states across a vast family of overlapping subsystems, many "pathological" alternative generators are excluded. Moreover, if we impose locality and symmetry as compression constraints, the space of candidate generators shrinks dramatically. In practice, identifiability is

recovered not because the trajectory alone forces it, but because the simplest consistent generator that explains all subsystem behaviors is sharply constrained.

Still, there may remain multiple viable generators even after all reasonable constraints. When that happens, the correct conclusion is not "we need more data" (we already have all the data), but "the theory is underdetermined by the history." At that point the best one can do is characterize the family of generators consistent with the dataset, identify which features are invariant across the family (hence genuinely deduced), and acknowledge which features depend on interpretive or representational choices.

This is where the paper's overall discipline matters. The goal is not to claim that omniscience erases ambiguity. The goal is to map where ambiguity truly survives and why. In the reconstruction of dynamics, that ambiguity takes a precise form: the universal state history constrains the generator strongly, but it may determine only an equivalence class unless additional principles—locality, symmetry, simplicity, and subsystem stability—select a preferred representative.

4. Deciding Unitarity: Collapse, Noise, or Pure Decoherence

Once the universal state history is available, the most emotionally charged question in quantum foundations becomes an engineering question: does the universe evolve unitarily, or not? In the ordinary scientific situation we infer this indirectly from measurement statistics, interference experiments, and consistency arguments. Here we are given the entire state history itself—$t \mapsto \rho(t)$, and in the strongest case $t \mapsto |\Psi(t)\rangle$. That changes the epistemic landscape. Unitarity is no longer a metaphysical posture; it becomes a property of the dataset that can be tested.

But "tested" does not mean trivial. Three confusions must be avoided.

First, decoherence is often mistaken for nonunitarity. Decoherence is fully compatible with global unitarity: locally things look mixed because they are entangled with environments. Second, noise can be objective (a genuine nonunitary law) or effective (coarse-grained ignorance). Third, the "Born rule" is usually introduced as a rule for assigning probabilities to measurement outcomes. If the universal state is known, what does it mean to speak of probability at all?

This section offers a disciplined way to separate these issues. We build criteria for (i) global purity and reversibility, (ii) objective collapse versus mere entanglement-generated mixing, (iii) model-based parameter inference if collapse is present, and (iv) a re-interpretation of the Born rule as a derived constraint on typical observers and records rather than a primitive instruction manual.

4.1 Global purity tests and reversibility as empirical criteria

If we are given the universal density operator ρ(t) for all t, the most direct diagnostic of unitarity is global purity.

Define purity:

$P(t) = Tr[\rho(t)^2]$

If the universe is in a pure state, P(t) = 1. If it is mixed, P(t) < 1. In an exactly unitary evolution of a closed system, purity is preserved. More generally, the entire spectrum of ρ(t) is preserved under unitary conjugation:

$\rho(t_2) = U \rho(t_1) U^\dagger$

which means all eigenvalues of ρ remain constant in time. So there are two related tests:

- **Purity preservation** (P(t) constant, and equals 1 if initially pure).
- **Spectral invariance** (eigenvalue distribution of ρ(t) constant).

In ordinary physics these tests are unreachable because we cannot access the universal ρ(t). In the thought experiment they are straightforward computations.

Reversibility provides a complementary criterion. If evolution is unitary, then there exists an inverse map:

$\rho(t_1) = U(t_1,t_2) \rho(t_2) U(t_1,t_2)^\dagger$

and the dynamics is, in principle, reversible. In nonunitary dynamics described by a general quantum channel E(t2,t1), an inverse channel may not exist, and even when a left-inverse exists on a restricted subset of states, it generally fails globally. So another test is:

- **Existence of a consistent inverse evolution**: does there exist a family of maps that reverses the forward evolution for all times while remaining physically admissible?

With complete access, reversibility can be evaluated without appealing to “measurement collapse.” If ρ(t) exhibits abrupt discontinuities, spectral changes, or irrecoverable loss of information at the global level, those are signatures of nonunitarity.

However, one must be careful. A globally pure state can still look locally irreversible because entanglement spreads. From the standpoint of a subsystem, information seems to be “lost” into the environment. But in a globally unitary universe, that information is not destroyed; it is redistributed into correlations. This is exactly why global tests are decisive in the thought experiment. The question is not whether subsystems thermalize, but whether information is conserved in the total state.

So the first conclusion is sharp: with full ρ(t), the unitary versus nonunitary distinction becomes an empirical property of the dataset, visible in purity and spectrum preservation and in the existence (or impossibility) of global reversal.

4.2 Distinguishing environmental decoherence from objective collapse

Decoherence is the mundane workhorse that makes the macroscopic world look classical. It is also the most common source of conceptual confusion, because it produces phenomena that resemble collapse while remaining fully unitary at the global level.

In a globally unitary universe, the universal state $|\Psi(t)\rangle$ evolves smoothly. A subsystem A, however, generally has a reduced state:

$$\rho_A(t) = \mathrm{Tr}_{A^c}[|\Psi(t)\rangle\langle\Psi(t)|]$$

which is typically mixed. Decoherence occurs when interaction with the environment produces rapid suppression of interference between certain components of A in a preferred basis (a "pointer basis"). In reduced-state language, off-diagonal terms in that basis become small. In everyday lab language, interference disappears and outcomes become stable.

Objective collapse theories claim something stronger: that collapse is not merely an effective description of entanglement, but a real, dynamical reduction process. Collapse changes the global state in a way that is not unitary and not reversible. It is a change in the laws, not a change in bookkeeping.

With complete access, we can separate these cleanly by asking whether apparent collapses in subsystems correspond to genuine global nonunitarity.

A principled distinction can be stated in three diagnostic questions:

(i) Is global purity preserved?
If global purity remains exactly 1 and the global state evolves continuously under a unitary operator, then there is no objective collapse. Whatever "collapse-like" behavior we see in subsystems is decoherence and record formation.

(ii) Are "branch weights" stable under unitary evolution, or do they show stochastic kicks?
In unitary decoherence, branch structure emerges as a decomposition of $|\Psi(t)\rangle$ into effectively noninterfering components relative to a chosen coarse-graining. Those components and their squared amplitudes evolve deterministically (given the Hamiltonian). In collapse dynamics, there are typically stochastic elements that modify weights in a way not representable as unitary evolution. Full access allows us to check whether the global amplitude structure evolves smoothly or exhibits random discontinuities.

(iii) Is information destroyed or merely dispersed?
Under unitary decoherence, information about phase relationships is not destroyed; it becomes encoded in highly nonlocal correlations. Under objective collapse, some information is genuinely erased. With full state access, we can quantify whether correlations conserve the full information content or whether there is irrecoverable loss.

This is where the thought experiment has its strongest clarifying power. In ordinary discussions, "collapse" is invoked partly because we cannot track the environment's degrees of freedom. In a full-state dataset, that excuse disappears. If the universe is unitary, then what looks like collapse to internal observers is explained entirely by entanglement structure and the emergence of stable records. If the universe is not unitary, the nonunitarity is visible in the global object itself, not merely in our narratives.

So "decoherence versus collapse" becomes a direct data question: does the full ρ(t) show signatures of nonunitary evolution, or does all apparent reduction occur only at the level of reduced states via tracing out degrees of freedom?

4.3 Parameter inference for collapse models: rates, length scales, and couplings

Suppose the global tests do indicate nonunitary dynamics consistent with objective collapse or fundamental noise. With complete ρ(t), we can go beyond yes/no and infer parameters.

Collapse models typically introduce additional terms into the evolution law that cause spatial or mass-dependent localization, with specific rates and length scales. The details vary by model, but the inferential structure is common: nonunitarity should leave systematic fingerprints in how coherence is lost as a function of system size, mass distribution, separation in configuration space, and time.

With full access, parameter inference would proceed conceptually as follows.

(i) Identify the form of deviation from unitary evolution
From ρ(t) we infer a best-fitting generator L_t such that $d\rho/dt = L_t(\rho)$. We then separate L_t into a unitary-like component (commutator with an inferred Hamiltonian) plus additional terms responsible for collapse/noise. The simplest decomposition is not guaranteed unique, but the goal is to find a minimal representation of the deviation.

(ii) Extract scaling laws from coherence decay
Pick families of subsystems and track how their reduced-state coherence decays in relevant bases. In unitary decoherence, decay rates depend on coupling to environments. In objective collapse, there are predicted universal scalings (e.g., stronger localization for larger masses or

larger spatial separations). With full state access, one can compute coherence measures directly and fit how they depend on mass, geometry, separation, and time.

(iii) Infer characteristic length and time scales
Collapse theories often posit a localization length scale and a collapse rate. In the dataset these would appear as characteristic scales at which superpositions become unstable even in the absence of environmental entanglement explanations. Because we can condition on the full environment, we can ask: after factoring out unitary entanglement effects, is there still an intrinsic decohering tendency with a fixed scale? If yes, its parameters are inferable.

(iv) Identify coupling structure
Nonunitary terms might couple to mass density, particle number, energy, or other operators. Complete access allows us to test which operators best predict where and how collapse-like effects occur. The "best compression" principle returns here: choose the coupling structure that most compactly accounts for the pattern of nonunitary deviations across all subsystems.

Two cautions matter.

First, nonunitarity can be mimicked by coarse-graining and by choosing a poor subsystem factorization. A careful analysis must show that inferred collapse parameters are robust under reasonable refactorizations, or else acknowledge that the "collapse-like generator" is representation-dependent.

Second, even with perfect data, identifiability can fail if the universe's actual trajectory does not sample enough situations to distinguish two parameter sets that would differ only in counterfactual conditions. However, the universe is large and contains an enormous variety of physical contexts. In practice, a complete universal history is likely to include abundant regimes—microscopic interference, mesoscopic systems, macroscopic record formation—providing strong leverage to infer scaling behavior.

So if objective collapse is real, full-state access would not merely settle the debate; it would let us reverse-engineer the collapse law quantitatively.

4.4 What "Born rule" becomes when the universal state is fully known

The Born rule is usually introduced as: the probability of an outcome equals the squared amplitude of the corresponding component of the wavefunction. In textbook practice, it is an algorithm for predicting measurement statistics.

But in our thought experiment, the universal state is not merely a tool for prediction; it is part of the dataset. We can see the actual history of branch formation, record creation, and the correlations that internal observers call "outcomes." So what role is left for the Born rule?

There are two distinct roles it can play, depending on what is true about unitarity.

If the universe is unitary:
Then probabilities are not fundamental at the global level. The universal evolution is deterministic (given the Hamiltonian and initial state), and the state contains all branches. Probability becomes an emergent, observer-relative concept tied to self-location within branching structure.

In that setting, the Born rule becomes a constraint on rational expectation for internal observers who (i) find themselves in a world with stable classical records and (ii) do not have access to global phase information across branches. The Born weights then function as the only stable measure that makes sense of typicality across branches and ensures consistent long-run frequencies in repeated experiments as encoded in records. In other words: the Born rule is no longer "a law of collapse," but a law governing the statistics of records across decohered branches under unitary dynamics.

With full access, one could verify this directly: take repeated experimental structures embedded in the universal history and check whether the distribution of branch records aligns with squared-amplitude weights. The rule becomes empirically checkable as a statement about the measure induced by $|\Psi(t)\rangle$ on the ensemble of record-bearing branches.

If the universe is nonunitary (objective collapse/noise):
Then probabilities may be fundamental in the dynamics itself. Collapse events, if genuinely stochastic, introduce irreducible randomness. In that case, the Born rule can appear as a structural property of the collapse process: the stochastic law that governs how outcomes are realized. Full-state access would let us distinguish whether Born weighting is exact, approximate, or modified, and under what conditions.

In both cases, the Born rule is demoted from a mysterious instruction manual for "measurement" to a mathematically testable regularity of the universal history:

- Under unitary evolution: a regularity about branch measures and record statistics.
- Under objective collapse: a regularity about the stochastic kernel of collapse events.

Either way, the key conceptual shift is this: when $|\Psi(t)\rangle$ is known, "probability" is not about ignorance of the state. It is about the structure of records and the perspective of embedded observers. The Born rule becomes a bridge principle connecting universal amplitude structure to the experienced frequencies encoded in stable classical memory.

This is one of the thought experiment's cleanest dividends. It replaces the vague question "why do we see definite outcomes?" with a precise question: what dynamical and informational conditions on ρ(t) yield stable record structures, and what weighting on those record structures matches the long-run statistics found within the history? Under complete access, the answer is not philosophized; it is computed.

5. Symmetries and Conservation Laws as Algorithmic Invariants

In a universe where we possess the full quantum state history, symmetry stops being a metaphysical guess and becomes an informational fact. Symmetries are precisely those transformations under which the description can be made shorter: they are redundancies, repetitions, and invariances that compress the generator of evolution. In classical physics, symmetry is inferred from limited experiments; in the present thought experiment, symmetry is inferred from the inferred law itself and verified against the entire trajectory.

This section develops that viewpoint. If we have reconstructed a generator of time evolution—unitary Hamiltonian dynamics, a more general channel generator, or some deeper law—then symmetries appear as invariances of that generator. Conservation laws then follow as algorithmic invariants: quantities whose values are fixed along the entire history, not because we declared them conserved, but because the state history and the generator jointly imply that no consistent evolution can change them.

However, "symmetry" is not one thing. Some symmetries are physical; some are gauge redundancies—different descriptions of the same underlying state. Complete access to $|\Psi(t)\rangle$ can reveal many invariances, but it cannot always tell us whether an invariance is ontic (a real symmetry of nature) or representational (a redundancy of our chosen variables). The distinction matters: physical symmetry tells you something about the world; gauge redundancy tells you something about how you wrote the world.

We proceed in four steps: (i) recovering Noether structure from the generator, (ii) separating gauge from physical symmetry and stating what the dataset can fix, (iii) detecting symmetry breaking and phase transitions as changes in the organization of $|\Psi(t)\rangle$, and (iv) inferring hidden sectors through missing conserved charges and dynamical residuals.

5.1 Noether structure recovered from the inferred generator

In its cleanest form, Noether's theorem links continuous symmetries to conserved quantities. If the dynamics is generated by a Hamiltonian H and the system has a continuous symmetry generated by an operator G, then invariance implies:

$[H, G] = 0$

and the expectation of G is conserved in time (in the appropriate conditions). In a broader setting, symmetries correspond to transformations that commute with the dynamics.

In our thought experiment, we are not handed H by divine fiat; we infer it as the best generator that compresses $|\Psi(t)\rangle$. Once H (or its nonunitary analogue) is inferred, we can algorithmically recover its symmetry structure:

- Search for groups of transformations U_g (unitaries representing symmetry actions) such that H is invariant under conjugation: $U_g H U_g^\dagger = H$.
- Equivalently, search for generators G such that $[H, G] = 0$, allowing for time dependence if H(t) varies.

This is "Noether recovery" as a computational problem: identify the maximal symmetry group consistent with the inferred generator, and then identify the associated conserved quantities.

But the point is not merely to recover textbook symmetries. With complete access, we can go further:

- We can detect approximate symmetries: operators G for which [H, G] is small in a quantified sense, implying nearly conserved quantities and slow drifts.
- We can detect emergent symmetries: symmetries not present at microscopic scale but arising in effective descriptions after coarse-graining (for example, scale invariance near critical points).
- We can detect broken symmetries: cases where the generator is symmetric but the state history selects a particular sector (spontaneous symmetry breaking), or where the generator itself changes regime so that symmetry is lost (explicit breaking).

From the algorithmic-invariant viewpoint, conservation laws are not primarily "laws about numbers." They are constraints on the compressibility of the evolution. If a quantity is conserved, then the history can be segmented into invariant sectors labeled by its value, reducing the complexity of describing transitions (because transitions are forbidden). Conservation is a compression: it deletes whole families of possible trajectories.

With full $|\Psi(t)\rangle$, we can verify conservation at two levels:

- At the generator level: the commutation/invariance relations.
- At the history level: the invariance of expectation values, distributions over eigenvalues in sectors, and the impossibility of trajectories crossing between superselection-like sectors.

In other words, Noether's theorem becomes not merely a principle we trust, but a structural relationship we can rediscover and validate from the complete dataset.

5.2 Gauge redundancy versus physical symmetry: what the data can and cannot fix

A central subtlety is that not all invariances are created equal. Some transformations represent physically distinct possibilities that leave the laws invariant—spatial translations, rotations, time translations, internal global symmetries. Others represent changes of description that do not correspond to physically distinct states—gauge redundancies.

With complete access to $|\Psi(t)\rangle$, can we tell the difference?

We can detect invariances of the inferred generator, but distinguishing gauge from physical symmetry requires asking: does the transformation map physically distinct states to distinct states, or does it map one description of a state to another description of the same state?

In many frameworks, gauge transformations act trivially on physical observables; they alter unphysical variables while leaving all gauge-invariant quantities unchanged. This implies an important limitation: if two descriptions are exactly gauge-equivalent, then no dataset of physical states and correlations can distinguish them, because they represent the same physical content.

That means "complete access" does not magically reveal gauge choice as a fact of nature. It reveals invariants. The dataset can force us toward gauge-invariant structures because only those are stable across equivalent descriptions. But it cannot, by itself, elevate a gauge convention to an ontological truth.

Practically, this yields a disciplined conclusion:

- The dataset can identify the algebra of gauge-invariant observables and show that dynamics closes on that algebra.
- The dataset can reveal constraints (like Gauss-law-type relations) that indicate redundancy in the underlying variables.
- The dataset can show that certain apparent degrees of freedom never carry independent information across the history—an algorithmic signature of redundancy.

But the dataset cannot uniquely determine "the gauge potential" as a physical object if only gauge-invariant combinations are observable and dynamically meaningful. Likewise, if two theoretical formulations are related by a duality that preserves all invariant content, the complete history may not select one formulation over the other.

So we adopt a rule that will recur throughout the paper: treat gauge-dependent objects as representational, and treat gauge-invariant objects as deduced. The strongest claims of inference should be stated in terms of invariants under the equivalence transformations that leave the dynamics and all record-structures unchanged.

This matters not only for internal consistency but also for ontology. When people say “the world is made of fields” or “the world is made of particles,” they are often choosing a representation. The dataset can privilege a representation by compressibility and locality, but when representations are truly equivalent, physics properly lives in the shared invariants.

5.3 Detecting symmetry breaking and phase transitions from $|\Psi(t)\rangle$

Symmetry breaking and phase transitions are, in a sense, the moments where the universe changes its own grammar. The microscopic laws might remain the same, but the effective degrees of freedom and the stable patterns of behavior reorganize. With full access to $|\Psi(t)\rangle$, we can detect these reorganizations directly—not merely via macroscopic observables, but via changes in correlation structure, entanglement patterns, and the stability of effective descriptions.

There are two main categories to distinguish.

Spontaneous symmetry breaking (SSB)
In SSB, the generator remains symmetric, but the state (and its record structure) selects a particular sector among many symmetry-related possibilities. A classic signature is that the state’s long-range order picks a direction or phase, making certain expectation values nonzero even though the Hamiltonian is invariant.

In our dataset-based approach, SSB can be detected by:

- Identifying symmetry of the generator (invariance group of H).
- Observing that the state history occupies a submanifold not invariant under the full symmetry group, but mapped into other equally valid submanifolds by the symmetry action.
- Detecting robust long-range correlations and order parameters that remain stable over time once the symmetry is broken.

Importantly, with complete access we can distinguish true SSB from mere choice of coordinates. The broken symmetry should be reflected in invariant patterns: correlation functions that exhibit long-range order, sector selection that persists under perturbations, and the emergence of superselection-like behavior in the effective description.

Explicit symmetry breaking and dynamical regime change
Sometimes the generator itself changes form across regimes: effective parameters vary, new terms become important, or a deeper theory yields different effective symmetries at different scales. In cosmology one thinks of symmetry breaking in the early universe, but the principle is general: when the best compression of the history requires different effective generators in different time windows, we can interpret that as a dynamical phase structure.

Phase transitions, in this broad sense, can be detected by:

- Changes in the optimal effective variables: a new coarse-graining suddenly becomes far more compressive.
- Changes in entanglement scaling and correlation length: critical phenomena often show distinctive scaling relations.
- Appearance or disappearance of emergent symmetries: near critical points, scale invariance and conformal-like behavior can emerge even if absent microscopically.

Because $|\Psi(t)\rangle$ is fully known, phase transitions are not just changes in thermodynamic variables. They are changes in the geometry of the state within Hilbert space: how correlations organize, how entanglement distributes across partitions, and which descriptions become stable.

One practical conceptual benefit follows. Much of physics is built around detecting phase transitions from limited macroscopic probes. Here, we can define them more fundamentally: a phase transition is a reconfiguration of the state history such that the minimal effective generator changes class. That is an algorithmic criterion, not an anthropocentric one.

5.4 Hidden sectors: inferring conserved quantities tied to unknown degrees of freedom

"Hidden sector" is a modern phrase, but the inferential logic is old: when what you can track does not close under the dynamics, something else is there.

In the complete access setting, this becomes striking. If we truly have the universal state $|\Psi(t)\rangle$, then nothing is hidden from us in principle. But something can still be "hidden" in the sense that our chosen decomposition, our candidate generator class, or our assumed degrees of freedom are incomplete. The dataset may force us to enlarge the model.

Hidden sectors can appear through at least three kinds of signatures.

(i) Conservation without visible carriers
Suppose we infer that some operator G is conserved (it commutes with the generator) but we

cannot express G in terms of the degrees of freedom we thought were fundamental in our representation. That is an alarm bell. It suggests that the true minimal description includes additional degrees of freedom or a different factorization in which G becomes natural.

Historically, this is how new entities are discovered: symmetry implies charge; charge implies carriers. If the charge is conserved but carriers are absent in our ontology, our ontology is incomplete.

(ii) Dynamical residuals that resist compression
When fitting a generator, we may find systematic residuals: patterns in the evolution that cannot be captured by any local/symmetric/simple law in the assumed degrees of freedom. Introducing an additional sector—new fields, new modes, new constraints—may drastically improve compression. In this sense, hidden sectors are discovered as the minimal extension that makes the generator simple.

(iii) Entanglement structure inconsistent with the assumed degrees of freedom
If we choose a factorization into subsystems and find entanglement patterns that are "too structured" or "too persistent" to be explained by interactions in that factorization, it may indicate that our factorization is misaligned. Re-factorizing the Hilbert space can reveal latent separations or couplings. Alternatively, it can reveal additional degrees of freedom required to represent the correlations locally.

In practical terms, the analysis becomes: find the representation in which the generator is maximally local and symmetric and in which conserved charges have natural carriers. If achieving this requires adding degrees of freedom, that is evidence for hidden sectors.

This is not restricted to particle physics. One can think of "hidden sectors" as any additional structure required to close the dynamics and compress the history: gravitational degrees of freedom if you began with matter-only variables; topological sectors if local fields are insufficient; or even emergent variables (like quasiparticles) that are not fundamental but are the right carriers at a given scale.

The underlying epistemic point is consistent with the paper's theme: conservation laws and symmetries are algorithmic invariants. When invariants appear that cannot be expressed cleanly in a given ontology, the ontology is being outcompeted by a better compression. Hidden sectors are the universe's way of telling the analyst: your variable set is not the right one yet.

In summary, complete access to $|\Psi(t)\rangle$ does not merely allow us to test known symmetries; it turns symmetry discovery into a principled inference problem. Noether structure is recovered from the inferred generator, gauge redundancies are recognized as representational

equivalences that the data cannot break, symmetry breaking and phase transitions are detected as reorganizations of correlation and compressibility structure, and hidden sectors are inferred when invariants and residuals demand a richer set of degrees of freedom. In this setting, "symmetry" is not an aesthetic preference. It is the skeleton of compressibility that the universal history itself imposes.

6. Locality and Causality: Building the Universe's Influence Graph

In ordinary physics, locality and causality are not merely scientific principles; they are the scaffolding of intelligibility. We assume influences propagate with finite speed, that causes precede effects, and that distant correlations do not automatically imply a direct causal link. But these assumptions are historically justified by sparse experiments, not by omniscient access to the full state. In the present thought experiment, locality and causality become inferable structural properties of the universal history.

The central goal of this section is to show how, given complete access to $t \mapsto \rho(t)$ (or $t \mapsto |\Psi(t)\rangle$), one could reconstruct an "influence graph" of the universe: a representation of which degrees of freedom can affect which others, through what pathways, with what characteristic delays, and under what constraints. This graph is not merely a schematic picture; it is the causal analogue of the inferred dynamical generator. If the generator tells us "what evolves," the influence graph tells us "how influence moves."

Quantum mechanics complicates causal inference because correlation is cheap. Entanglement creates strong correlations that cannot be explained by classical common causes, yet those correlations do not allow superluminal signaling. Meanwhile, the very idea of intervention—doing X to see if Y changes—is delicate in a closed universe. If we have the entire universal history, can we do causal discovery without interventions?

The answer is: not in the strictest experimental sense, but yes in an inferential sense. With complete state access and a reconstructed generator, we can infer conditional dependencies, propagation constraints, and effective causal horizons. We can distinguish the correlation structure induced by entanglement from the directional influence structure induced by dynamics. And we can state what would count as a genuine violation of locality: not merely Bell correlations, but the presence of controllable, propagating influence outside any finite-speed bound, robust under representation choices.

6.1 Causal discovery in quantum systems: interventions without interventions

Causal discovery traditionally leans on interventions: you set a variable and see what changes. In a closed universe, “intervene” is a tricky word. There is nothing outside the universe to set its variables. Even an internal agent is just part of the evolving state, and any “choice” is itself a physical process correlated with the rest of the world.

So how can we speak about causality at all?

The key move is to replace literal interventions with two inferential substitutes that become available under complete access:

(i) Counterfactual simulation from an inferred generator
If we infer a dynamical generator—Hamiltonian or more general—then we can consider alternative initial conditions (or alternative micro-perturbations) and evolve them forward in simulation. These are not interventions in the realized history; they are lawful counterfactuals consistent with the inferred dynamics. The causal question becomes: if we slightly vary subsystem A at time t, while holding the rest fixed as much as the formalism allows, what changes appear in subsystem B at later times? This defines influence operationally as sensitivity under dynamical variation.

Complete access strengthens this because we can define “holding the rest fixed” in a principled way: we can work at the level of the universal state, perturb a localized component, and propagate forward using the inferred generator. In ordinary science, we rarely know enough to set up that kind of precise counterfactual. Here, we do.

(ii) Exploiting natural experiments embedded in the universal history
Even without counterfactual simulation, the universe’s history contains an immense variety of repeated motifs: similar subsystems prepared in similar states under different surrounding conditions. These are “natural interventions” in the sense that the world itself supplies variation. With full access, we can search the history for matched contexts and treat them as quasi-experiments: identify episodes where subsystem A differed while relevant confounders were similar, then examine the downstream differences in B.

This is causal inference by pattern mining rather than by laboratory control. Under complete access, it becomes extremely powerful because the dataset is not a handful of trials; it is the entire universe across all times and locations. The limiting factor is not sample size; it is defining equivalence classes of “similar contexts” in a principled quantum way.

These two tools allow a careful redefinition:

- A causal influence relation A → B exists if perturbations or natural variations in A at time t systematically produce changes in B at later times, in a way that is stable under coarse-graining and not explainable solely as common-cause correlation.

Notice the humility embedded here. Causality is not treated as metaphysically primitive; it is treated as a stable asymmetry in the inferred dynamics and in the sensitivity structure of the state history.

6.2 Light cones, finite propagation, and effective causal speed limits

Locality, in the most operational sense, is the claim that influence propagation is bounded: there exists a maximum speed at which information or disturbance can travel. In relativistic physics, that bound is set by the light cone. In other frameworks, an analogous causal horizon might exist.

With complete access, we can infer such speed limits as properties of the influence graph and the generator.

The idea is conceptually simple: if we perturb a localized subsystem A at time t, and we then examine where any detectable change appears later, we can map the boundary of influence. If changes in distant subsystems appear only after delays consistent with finite propagation, then the universe has an effective speed limit. If changes appear instantaneously at arbitrary distance in a controllable way, then it does not.

But to do this, we need a notion of "distance" and "locality" in the first place. That seems circular: we want to infer spacetime structure, yet we're referencing it.

The way out is to proceed in two stages.

Stage 1: Infer adjacency from dynamical coupling
Given a candidate factorization of the Hilbert space, infer which degrees of freedom are directly coupled in the generator. If the generator can be expressed as a sum of terms each acting on small neighborhoods, those neighborhoods define an adjacency relation. This adjacency can then be treated as the seed of an emergent geometry: degrees of freedom that couple strongly and directly are "near," those that couple only through many intermediates are "far."

Stage 2: Infer propagation speed in the emergent adjacency metric
Once we have an adjacency structure, we can study how perturbations spread across it. The influence graph then yields an effective "cone": a bound on how the support of a localized

perturbation grows over time. This is the causal analogue of a light cone, defined without presupposing spacetime coordinates.

If the universe already admits a relativistic spacetime description, this recovered structure should align with ordinary light cones and the speed of light. If it does not, we will still recover some effective propagation constraint if the dynamics is local in any meaningful sense.

Importantly, "finite propagation" need not mean strict zero influence outside the cone. In quantum systems, there can be exponentially small tails depending on definitions and norms. The operational content is: there is a sharp, representation-stable bound such that influence outside the cone cannot be amplified into controllable signaling faster than the bound. Under complete access, we can quantify this precisely in terms of how distinguishable reduced states become under localized perturbations as a function of time and graph distance.

So light cones become a derived object: a boundary in influence space, not an axiom.

6.3 Separating entanglement correlations from causal influence

This is the point where many discussions derail. Entanglement creates correlations that look "nonlocal," and Bell's theorem rules out classical local hidden-variable explanations. But Bell correlations do not imply faster-than-light causal influence. They show that the joint statistics cannot be explained by classical local realism; they do not show that one party can signal to the other by choosing a measurement.

With complete access to the universal state, we can separate these two notions cleanly.

A correlation is a pattern in joint statistics or mutual information. Causal influence is about controllable change: does varying A produce a change in B's reduced state, accessible to B, without conditioning on information from A?

In formal terms (without leaning on any particular measurement narrative), we can use reduced states and their dependence:

- If changes in A do not change ρ_B (the reduced state of B), then A cannot signal to B, regardless of entanglement. Entanglement may change correlations between A and B, but B's local state remains unchanged.
- If changes in A do change ρ_B in a way that is accessible to B locally, then A influences B.

In ordinary quantum mechanics under unitary evolution and local operations, entanglement alone does not allow signaling, because local operations on A do not change the reduced state

of B in a way that transmits controllable information. This is the familiar no-signaling structure, but here it is not asserted; it is checked directly from the universal dynamics.

Complete access gives us three concrete separation tests:

(i) Local-state invariance tests
Identify entangled subsystems A and B. Vary A's internal configuration in controlled simulations (or find natural variants), then compute whether ρ_B changes. If it does not, the correlation is not a direct causal influence channel.

(ii) Conditional versus unconditional dependence
Even if B's state depends on A when conditioned on additional information (for example, knowing a measurement outcome), that does not imply direct influence. Influence requires unconditional change accessible to B without extra side information. With full state history, we can compute both conditional and unconditional dependencies and separate them cleanly.

(iii) Mediator tracing
If A and B correlations evolve, we can search for mediating degrees of freedom M that carry the dynamical influence between them. In a local theory, any apparent A–B dependence should be explainable by interactions propagating through mediators along adjacency pathways. Entanglement can be created by local interactions and then persist without ongoing influence. The influence graph tells us where the mediators are and whether the correlations require direct nonlocal coupling.

This separation is critical for diagnosing genuine nonlocal dynamics. Bell-type nonclassicality is compatible with locality of influence in the relativistic sense. The violation we are looking for in this paper is not "quantum weirdness," but breakdown of finite-speed influence.

6.4 Diagnosing nonlocal dynamics: what would count as a genuine violation

If the dataset reveals correlations, that is not enough. To claim genuine nonlocal dynamics—something beyond unitary local quantum mechanics—we need a criterion that is robust, operational, and not an artifact of representation.

A genuine violation would require at least one of the following, stated in dataset terms.

(i) Superluminal signaling in the influence graph
There exists a decomposition in which A can be perturbed locally and B's reduced state changes in a way that is accessible at B, outside any finite propagation bound, and this effect cannot be removed by changing factorization in any reasonable "locality-favoring" way. In other words,

there is a stable, compressively preferred description in which controllable influence outruns the recovered causal cone.

This is stronger than Bell nonlocality. It is causal nonlocality.

(ii) A generator that cannot be localized without losing compression
Suppose we try every plausible factorization and coarse-graining and find that any generator capable of reproducing ρ(t) must include long-range couplings between distant degrees of freedom. If these couplings are not removable by gauge or duality transformations and if they persist as the minimal description, then nonlocal dynamics is not a choice; it is a demanded feature.

In algorithmic terms: locality fails as a compression principle. The simplest generator is inherently nonlocal.

(iii) Violations of causal composition consistent with finite-speed propagation
Local dynamics typically yields a structured causal ordering: influences propagate through intermediate regions, producing constraints analogous to domain-of-dependence properties. If the dataset forces violations of such constraints—events whose dependence structure cannot be mediated through any intermediate region in any compressive representation—that is evidence of nonlocal influence.

(iv) Robust non-Markovian "action at a distance" that resists mediator modeling
Some apparent nonlocal effects can be explained by hidden mediators or by misidentified subsystems. A genuine violation would be an influence pattern that cannot be captured by any model with mediators, even allowing for complex environments, without introducing explicit nonlocal coupling terms. Again, complete access makes this test sharp: we can include mediators explicitly, because nothing is hidden.

A final caution is essential. Nonlocality claims are notoriously sensitive to how "systems" are defined. Because factorization is not unique, a careless analysis can manufacture nonlocality by choosing a bad decomposition. That is why the paper insists on robustness and compression: the relevant nonlocality is what survives the search for a representation that makes the generator simplest and most local.

If, after that search, the best compression remains nonlocal, then nonlocal dynamics is not a philosophical provocation; it is an empirical property of the universe-as-data.

In summary, complete access to the universal state history allows us to reconstruct causality without external interventions by combining counterfactual simulation under an inferred generator with "natural experiments" mined from the universe's own variation. It allows us to

derive effective light-cone structures and causal speed limits from adjacency implied by the generator. It lets us separate entanglement correlations from causal influence using reduced-state dependence tests. And it provides stringent criteria for genuine nonlocal dynamics: not Bell correlations, but controllable influence outside any finite propagation bound, or an irreducibly nonlocal minimal generator. Under this thought experiment, locality and causality become computable features of the universe's structure—properties to be inferred, not presupposed.

7. Spacetime and Geometry: Fundamental Stage or Emergent Description

Spacetime is usually treated as the arena in which physics happens: events occur at points, fields live on manifolds, signals propagate along causal cones, and geometry tells matter how to move while matter tells geometry how to curve. But the thought experiment of complete quantum access invites a more radical inversion. If we are handed the universal state history directly, then "spacetime" is not automatically given; it is a candidate explanatory layer. The question becomes: is geometry a primitive part of the world's ontology, or is it a compressive representation extracted from patterns of dynamical coupling, correlation, and entanglement?

This section is not trying to settle quantum gravity. It is trying to clarify what the full-state dataset would and would not decide. In particular, if we can reconstruct a local influence graph (Section 6) and a generator with locality structure, we can treat that structure as the seed of an emergent geometry. Separately, if entanglement patterns encode adjacency and "distance," we can attempt to reconstruct geometry from information alone. These two routes—trajectory-based and entanglement-based—need not coincide. If they do, geometry earns the status of a privileged compression. If they conflict, geometry may be a useful but non-fundamental story, or our subsystem carving may be misaligned.

We develop four themes: two reconstruction routes, a criterion for when a metric is the best compression, what curvature/horizons/topology would look like in the complete record, and what could falsify the claim that spacetime is emergent.

7.1 Reconstructing geometry from trajectories versus from entanglement structure

There are two conceptually distinct ways to build "geometry" from the universal state history.

Route A: Trajectory-based geometry (dynamical adjacency)
If we have inferred a generator of evolution (Hamiltonian or channel generator), we can examine which degrees of freedom interact directly and strongly. This produces a coupling graph: nodes are degrees of freedom; edges represent direct terms in the generator. If the

generator is local in some representation, most edges connect only “nearby” nodes in that representation. Geometry can then be reconstructed as whatever embedding of this graph yields a coherent notion of neighborhood and distance.

This route treats “space” as the organization of interaction. It is fundamentally dynamical. Degrees of freedom are close if the generator couples them directly; far if coupling requires many intermediate steps. The resulting geometry is, in effect, a best-fit coordinate system for the generator’s sparsity pattern.

The strength of this route is its operational clarity: locality is defined by dynamical influence. The weakness is that it depends on factorization choices. A different decomposition of the Hilbert space can reshuffle what looks “local.” So the geometry recovered here must be judged by robustness and compression: does a particular emergent geometry make the generator dramatically simpler and more local than its competitors?

Route B: Entanglement-based geometry (information adjacency)
A second route begins not from the generator but from the structure of correlations and entanglement in $|\Psi(t)\rangle$. There is a broad family of ideas (with many variations) in which entanglement patterns define proximity: subsystems that share high mutual information or strong entanglement are “near,” while weakly correlated subsystems are “far.” In this view, geometry is a map that makes the correlation network look like a spatial manifold with decaying correlations with distance.

This route treats “space” as the organization of information. It is, in principle, static at each time slice: given the entanglement structure at time t, one can attempt to reconstruct a geometric picture without direct reference to the generator.

The strength of this route is that it uses data available directly from $\rho(t)$ even if the dynamical generator is hard to express locally. The weakness is that entanglement is not the same as causal adjacency. Entanglement can persist between systems that have not interacted for a long time. Correlations can be long-range without implying proximity. So entanglement-based geometry risks confusing “being correlated” with “being nearby.” This is precisely why Section 6 insisted on separating correlation from influence.

In the complete-access setting, we can compare these two geometries explicitly. If a geometry reconstructed from coupling structure matches a geometry reconstructed from entanglement structure across broad regimes, that concordance is evidence that geometry is not merely a visualization trick—it is a stable, cross-validated compression of both dynamics and information flow. If they disagree systematically, then at least one of the reconstructions is missing something, or geometry is not a single coherent concept at the fundamental level.

A subtle possibility is that both are right but at different scales or phases. In some regimes (say, near thermal equilibrium), entanglement structure might align with locality; in others (say, after long-range entanglement distribution), it might not. The dataset would allow these regime distinctions to be drawn sharply rather than argued abstractly.

7.2 When a metric is the best compression: information geometry criteria

A "metric" is not just a philosophical flourish. It is an extremely strong form of compression. If the universe admits a metric space description, then many structural facts become simple: "near things influence near things," correlations decay with distance, propagation respects cones, curvature encodes deviation from Euclidean expectations, and global properties constrain what can happen.

The paper's criterion for taking geometry seriously is therefore not "it feels intuitive," but "it compresses the generator and the history better than alternatives."

What does it mean, in this dataset-first setting, for a metric to be the best compression?

Criterion 1: Sparsity and locality of the generator in metric coordinates
A metric earns its keep if, when degrees of freedom are embedded in that metric space, the generator becomes local: coupling strengths are concentrated among near neighbors, and long-range couplings are negligible or structured in a simple way. If a metric embedding turns a dense, arbitrary interaction matrix into a sparse, approximately translation-invariant local interaction, that is decisive compression.

Criterion 2: Predictable decay of correlations with distance
Given a candidate metric, compute correlation functions or mutual information between regions as a function of distance. If correlations obey stable decay laws (exponential, power-law, or other simple scaling) in that metric, then the metric has explanatory power. If no distance notion makes correlations behave simply, then "distance" is not a natural organizing variable for that regime.

Criterion 3: Consistent causal cones
A metric and a causal structure should align: the recovered influence graph should yield propagation bounds that look like cones in the metric. If influence spreads in ways that are sharply bounded by distance-dependent delays, then geometry is doing work. If influence propagation is better explained by some non-metric structure (e.g., graph-theoretic constraints not embeddable in a low-dimensional manifold), then manifold geometry may be an overcompression.

Criterion 4: Low-dimensional manifold adequacy
A metric representation implies that the high-dimensional complexity of the degrees of freedom can be organized as points in a low-dimensional space (3+1 dimensions, or something close). This is a huge claim. With complete access, we can test manifold adequacy: does there exist a low-dimensional embedding that preserves the key adjacency and correlation structures with low distortion? If not, then the universe may be fundamentally high-dimensional in its relational structure, even if it appears 3D at macroscopic scales.

All of these criteria are "information-geometric" in the broad sense: they treat geometry as a tool for compressing relational data. A metric is accepted not because it is assumed, but because it is the coordinate system in which the dataset and the inferred generator become simplest.

This also clarifies what "spacetime is emergent" means operationally in this paper: it means that the metric and manifold are not primitive inputs but outputs of a compression procedure applied to $|\Psi(t)\rangle$ and its dynamics.

7.3 Curvature, horizons, and global topology in the full-state record

If we can reconstruct a metric-like structure, the next question is whether the full record allows us to infer curvature, horizons, and global topology—features that, in general relativity, are geometric and physically consequential.

In the dataset-first setting, these become patterns in how influence, correlation, and state organization behave.

Curvature as deviation from flat relational structure
Curvature, in geometric language, expresses how the space deviates from Euclidean expectations: triangle angle sums, geodesic divergence, volume growth with radius, and so on. In our setting, curvature corresponds to how adjacency and propagation patterns deviate from those of a flat metric embedding. For example:

- If neighborhoods grow faster or slower with radius than in flat space, that signals curvature-like structure.
- If geodesic paths (shortest influence pathways) show systematic focusing or defocusing, that is curvature in the influence graph sense.
- If correlation decay and propagation differ directionally, that signals anisotropic or inhomogeneous geometry.

With full access, we can compute these relational analogues directly by analyzing the influence graph and the best-fit metric embedding across time.

Horizons as information-access boundaries

Horizons are not just geometric surfaces; they are constraints on what can influence what, and what can be known by what. In the full record, a horizon is present if there exist regions such that influence from one region cannot reach another within the relevant time evolution, or such that information becomes effectively inaccessible.

Operationally, horizons appear as:

- persistent causal disconnection in the influence graph,
- irreversible scrambling where information about some region cannot be reconstructed from outside under the allowed dynamics,
- sharp boundaries in entanglement and correlation structure consistent with causal separation.

With complete access, we can distinguish true horizons from practical limitations. A true horizon is not "hard to see across," but "cannot influence across" under the generator.

Global topology as nonlocal consistency constraints

Topology is famously hard to infer from local data alone. But our dataset is not local: it is the total state history. Topology would manifest as global constraints on how local neighborhoods patch together. If the best-fit geometry is manifold-like, then topology shows up in:

- the presence of noncontractible loops in the adjacency structure,
- global identifications that force long-range correlations or constraint relationships,
- global modes or conserved quantities that correspond to topological sectors.

Even if the fundamental description is not a smooth manifold, the dataset may still imply topological invariants at an emergent level—features that are stable under local deformations and therefore act as robust compression targets.

The key point is that curvature, horizons, and topology become deducible only insofar as geometry itself has earned its keep as a compressive representation. If geometry is not the best compression, then these notions may be artifacts—useful metaphors rather than deduced structures. Complete access lets us decide that by testing whether the relational patterns are manifold-like or fundamentally graph-like or something else entirely.

7.4 What would falsify "spacetime is emergent" in this setting

The claim "spacetime is emergent" can be too elastic: if one is not careful, any outcome can be interpreted as supporting emergence. The thought experiment forces discipline by making falsification conditions explicit.

In this paper's setting, "spacetime is emergent" means: the best compression of the universal state history and its dynamics does not require assuming a spacetime manifold as primitive; instead, a spacetime-like description can be derived as a stable, low-dimensional organizing structure from the generator and/or entanglement patterns.

So what would falsify that?

Falsifier 1: Geometry must be input, not output, to achieve minimality
If no compression procedure starting from abstract $|\Psi(t)\rangle$ and generic dynamics can recover a spacetime structure without already presupposing geometric locality, then geometry is not emergent in the relevant sense. If the simplest generator cannot be stated without explicit spacetime primitives—points, neighborhoods, metric—then geometry is part of the irreducible law, not a derived convenience.

In other words: if assuming spacetime is necessary to write down the minimal generator, emergence fails as a compression claim.

Falsifier 2: No representation yields stable, low-dimensional locality
If, across all reasonable factorizations and embeddings, the generator remains irreducibly nonlocal or requires a high-dimensional relational structure that cannot be approximated by a low-dimensional manifold with low distortion, then spacetime-as-manifold is not an emergent description. One might still have an emergent graph geometry, but not spacetime in the usual sense. The "3+1D stage" would then be a parochial macroscopic approximation, not a deduced organizing structure.

Falsifier 3: Entanglement-based and dynamics-based reconstructions irreconcilably conflict
If geometry reconstructed from coupling/locality does not match geometry reconstructed from entanglement/information structure in any regime where both should apply, and if the mismatch is not explainable as a scale-dependent effect, then "geometry" may not be a coherent emergent object. That would suggest that spacetime is not a unifying derived structure but a patchwork of incompatible compressions.

Falsifier 4: Causal structure cannot be derived without presupposing it
If causal cones and speed limits cannot be inferred from the generator and state history without importing an external causal ordering (a background time or manifold structure), then

“emergent spacetime” loses content. In that case, causal structure would be a primitive input, and spacetime would not be derivable in the strong sense.

Notice what is not a falsifier: the mere fact that spacetime is a very good compression. Emergence is compatible with spacetime being extraordinarily effective—indeed, that is the point. What falsifies emergence is not effectiveness but irreducibility: geometry cannot be extracted without being assumed.

So the complete-access thought experiment gives a clean methodological stance. Either spacetime is the privileged output of compression from $|\Psi(t)\rangle$, in which case its metric, curvature, and horizons can be read off the influence and entanglement structures; or spacetime is a primitive coordinate system required to state the minimal generator, in which case it is fundamental at the level of law. The dataset does not force one conclusion by philosophy. It forces one conclusion by which description compresses the universal history most tightly while preserving its dynamical structure.

In summary, geometry can be reconstructed either from dynamical coupling (trajectory-based locality) or from entanglement/correlation structure (information adjacency). A metric earns its place when it provides the strongest compression: local generator sparsity, simple correlation decay, coherent causal cones, and low-dimensional manifold adequacy. Curvature, horizons, and topology become deduced patterns in influence and information constraints once geometry is justified. And “spacetime is emergent” is falsifiable here: it fails if geometry must be presupposed to state the minimal generator, if no low-dimensional locality representation exists, or if causal and informational reconstructions refuse to cohere.

8. Emergence of the Classical World: Records, Robustness, and Pointer Structures

The most stubborn feature of experience is not that the world is quantum, but that it feels classical. Tables have locations. Thermometers have readings. Memories persist. Different observers can compare notes and agree about what happened. In ordinary quantum foundations, this “classicality” is explained by a patchwork: decoherence suppresses interference, measurement selects outcomes, and classical probability emerges as an approximation.

With complete access to the universal state history, that patchwork can be replaced by something more precise: classicality is a pattern in the state’s information architecture. A “classical fact” is not a metaphysical primitive. It is a stable, redundantly encoded record across

subsystems, robust under perturbations and coarse-graining, such that a wide range of agents can reliably extract the same information without disturbing it.

This section formalizes that claim and shows what it becomes when we can compute everything—every reduced state, every mutual information, every entanglement entropy, every correlation function—exactly from ρ(t). Decoherence becomes bookkeeping, not mystery: it is the exact redistribution of phase coherence into correlations with the environment. Objectivity becomes redundancy: many disjoint fragments carry consistent information about the same variable. And the boundary of classicality becomes a compression boundary: classical descriptions work where they sharply reduce description length without losing predictive power, and fail where the state's structure refuses to be compressed into classical records.

8.1 What a "classical fact" is: stable records in subsystems

To say "a fact is classical" usually means two things: it is definite (not a superposition for all practical purposes), and it is shareable (multiple observers can learn it and agree). Both properties can be stated without invoking "collapse" if we treat the universe as a full quantum state and treat observers as subsystems.

Let A be a subsystem of interest (a pointer, a cat, a thermometer needle, a brain state), and let E denote an "environment" partition containing degrees of freedom that interact with A and can carry information away. In the universal state ρ(t), the classicality of some variable of A is tied to the existence of **records** of that variable in E.

A record is best defined operationally:

- A subsystem fragment F ⊂ E contains a record of A (or of some coarse-grained variable of A) if there exists a measurement on F that reveals information about that variable with high reliability, and doing so does not significantly disturb the record.

In density-operator language, this means the joint state ρ_{AF}(t) exhibits strong correlations in some preferred basis for A and corresponding distinguishable states of F. But "basis" must not be smuggled in. The preferred basis should be *inferred* as the one that makes records stable and robust.

This is where **pointer structures** enter. In decoherence theory, pointer states are those that remain stable under interaction with the environment: they get imprinted into E rather than being scrambled. In the full access setting, we can define pointer structure as:

- A set of states (or a coarse-grained observable) of A is pointer-like if, over time, correlations with E form in a way that preserves distinguishability among those states and suppresses interference between them, producing long-lived records.

Now define a “classical fact”:

- A classical fact at time t is a value of some pointer variable of A that is redundantly recorded across many disjoint fragments of E and remains stable over a timescale long compared to the relevant observation and interaction timescales.

This definition captures definiteness and shareability without metaphysical leaps. Definiteness corresponds to the reduced state ρ_A(t) being approximately diagonal in the pointer basis (or at least well-approximated by a mixture over pointer sectors). Shareability corresponds to redundant encoding: many fragments independently carry the same information.

Crucially, with complete access we can measure “stable” and “redundant” exactly, not as handwaving. Stability becomes a statement about how long the record correlations persist in ρ(t). Redundancy becomes a statement about how mutual information saturates across many fragments.

8.2 Decoherence as entanglement bookkeeping: exact mutual-information flows

Decoherence is often described as “the environment monitoring the system,” or as “phase information leaking into the environment.” These metaphors are useful, but vague. With full access to ρ(t), decoherence becomes an exact accounting of information flow and entanglement redistribution.

Consider the reduced state ρ_A(t). Even if the global state is pure and evolves unitarily, ρ_A(t) typically becomes mixed because A becomes entangled with E. The “loss of coherence” in A is not annihilation of information; it is transfer of coherence into correlations between A and E.

This can be formalized through information measures. Let S(X) denote the von Neumann entropy of subsystem X. In a globally pure state, entanglement entropy is bookkeeping: S(A) = S(E) for bipartitions A|E.

More finely, mutual information captures total correlation:

I(A:E) = S(A) + S(E) - S(AE)

In a globally pure state, S(AE)=0, so I(A:E)=2S(A). As decoherence proceeds, S(A) grows (until saturation), and thus I(A:E) grows, meaning A and E become more correlated.

But the real story of classicality is not just “A is correlated with E.” It is “many small pieces of E carry accessible information about A.” For that we consider fragments F of the environment and compute:

I(A:F)

As time evolves, I(A:F) typically rises as information about A gets imprinted into F. Under classic decoherence scenarios, there is a characteristic shape: even small fragments quickly acquire nearly the full classical information about A, and adding more of E yields diminishing returns. That shape is the fingerprint of redundant encoding.

With complete access, we can compute these curves exactly for any chosen partitioning of E into fragments. There is no need to talk about “for all practical purposes.” The mutual-information flow is literally available as data.

This framing has two clarifying consequences.

First, it shows precisely what decoherence is *not*: it is not a fundamental randomness injection and not a physical collapse. In unitary dynamics, decoherence is deterministic entanglement growth and correlation redistribution.

Second, it lets us distinguish decoherence from objective collapse in an unambiguous way (as in Section 4). If coherence is lost globally (purity changes), that is nonunitarity. If coherence is only lost locally (reduced states mix while the global state remains pure), that is decoherence-as-bookkeeping.

So decoherence becomes an exact thermodynamics of information: where phase relations go, how they spread, and which correlations become practically inaccessible to local observers.

8.3 Redundant encoding and objectivity: why many observers agree

The central puzzle of “classical reality” is not just why outcomes appear definite, but why they appear **objective**: why multiple observers can independently inspect the world and agree. The answer, in this dataset-first setting, is redundancy.

Suppose A is the system of interest and the environment E is partitioned into many disjoint fragments F_1, F_2, ..., F_n. Each fragment might be a chunk of scattered photons, air molecules, electronic fields, or any degrees of freedom that have interacted with A and carried away information.

Objectivity arises when:

- For many i, the fragment F_i contains nearly the full classical information about the pointer variable of A.
- Reading F_i does not appreciably disturb A or other fragments.
- Different fragments agree: the records are consistent.

This is why “many observers agree” is not a primitive axiom. It is a structural property of the universal state: the same information about A is copied (in a_toggle sense, not violating quantum no-cloning because it is classical information in a pointer basis) into many parts of E.

In mutual-information terms, redundancy can be quantified: how many disjoint fragments are sufficient to recover (almost) all the relevant information about A? In the full access setting, we can define an exact redundancy measure by searching over fragment sizes and counting how many fragments achieve a threshold fraction of the available classical information.

This has a powerful implication: objectivity is not universal. It depends on the existence of a pointer basis and on the environment’s ability to broadcast information about A. Many quantum variables do not get broadcast; their information remains nonlocally encoded, inaccessible to small fragments. Those variables are not classical facts.

So “agreement” among observers is explained by the environment acting as a communication medium that disseminates pointer information widely. The world appears classical because most macroscopic variables that matter—positions of large objects, pointer readings, macroscopic densities—are precisely those variables that naturally become redundantly recorded by ubiquitous interactions.

Complete access allows the argument to be made with no rhetorical handwaving. We can compute redundancy profiles across the entire universe’s history and identify where objectivity emerges and where it does not.

A final subtlety: redundancy depends on partitioning. If we choose a bizarre factorization of the Hilbert space, redundancy may vanish. That is not a refutation; it is a reminder of the paper’s recurring discipline. Classicality must be robust under reasonable decompositions. The “classical decomposition” is the one in which stable records form most naturally and which yields maximal compression of macroscopic behavior. In that sense, classicality helps select the right coarse-graining rather than being merely a product of it.

8.4 The boundary of classicality: where classical description stops compressing well

Classical physics is an extremely successful compression scheme. It replaces the full quantum state with a handful of variables—positions, momenta, fields, thermodynamic parameters—plus probabilistic noise where needed. But it is not universally valid. The full-state dataset lets us define the boundary of classicality in a crisp way: classical description stops being the best compression when the state's relevant structure cannot be captured by stable, redundantly recorded variables.

There are several ways this boundary appears.

(i) Lack of stable pointer variables
If interactions do not single out a preferred basis for a subsystem, then no variable becomes record-like. In such regimes, describing the subsystem classically is not compressive; you would need to track coherent superpositions and phase relations explicitly.

(ii) Absence of redundancy
Even if a subsystem decoheres (its reduced state becomes mixed), the information about its would-be classical variable might not be redundantly broadcast. If only a large, global portion of the environment contains the information, then small observers cannot access it, and "objectivity" fails. The variable is not a classical fact in the practical sense, and a classical model may mislead.

(iii) Strong quantum interference and contextual dependence
Classical models assume that probabilities are over definite states and that variables have joint distributions. Quantum structure violates this when noncommuting observables matter and interference terms materially affect outcomes. In full-state terms, the boundary is where off-diagonal structure and phase-sensitive correlations are not negligible in any robust coarse-graining.

(iv) High entanglement across the would-be classical cut
Classical descriptions typically assume weak entanglement between macroscopic degrees of freedom or that entanglement can be ignored except as noise. But in some regimes—strongly correlated many-body systems, near criticality, engineered quantum devices—entanglement is the main structure. In those regimes, classicality is a poor compression.

The paper's overarching thesis now becomes visible in a concrete case: the classical world emerges where it is the most efficient compression of the full quantum record. It does not emerge because someone collapses a wavefunction; it emerges because the universal dynamics produces stable, redundant records that make a classical narrative both accurate and compact.

And because we have the entire dataset, we can push the claim further. We can, in principle, map the universe into regions of "high classical compressibility" and "low classical compressibility," across time and scale. We can identify the precise thresholds where redundancy and stability appear, and track how those thresholds shift with environment coupling, temperature, density, and dynamical regime. Classicality becomes a phase of information architecture, not a philosophical afterthought.

In summary, a "classical fact" is a stable, redundantly encoded record in subsystems; decoherence is exact entanglement bookkeeping visible as mutual-information flow; objectivity arises because many disjoint environment fragments carry consistent records; and the boundary of classicality is the point where pointer structures and redundancy fail and a classical description no longer compresses the state history well. Under complete quantum access, the classical world is not assumed—it is detected as the universe's most successful lossy compression layer.

9. The Arrow of Time: Exact Accounting of Entropy and Irreversibility

Few facts feel as certain as time's direction. We remember the past, not the future. Ice melts, it does not un-melt. Signals disperse; neat arrangements become messy. Yet in much of fundamental physics—classical Hamiltonian mechanics and unitary quantum theory—the microscopic laws are time-reversal symmetric. The arrow of time is therefore not obviously in the laws; it is in the world's history.

That makes the arrow of time an ideal target for the full-state thought experiment. If we really have complete access to the universal state history, we can stop arguing about vague intuitions ("entropy tends to increase") and do exact bookkeeping: compute entanglement entropies, coarse-grained entropies, accessible mutual informations, and the irreversible formation of records, across all scales and times. We can then ask: what exactly increases? what stays constant? where does irreversibility arise in a globally reversible theory? and what boundary conditions, if any, are required?

This section builds an "entropy ledger" that aligns quantum information measures with thermodynamic entropy, and then uses that ledger to interpret the arrow as a statement about which correlations become accessible to internal observers and which become effectively lost in fine-grained entanglement. The guiding idea is consistent with the paper's central claim: the arrow is not a metaphysical primitive; it is a compressive regularity of the universal record.

9.1 Entanglement entropy versus thermodynamic entropy: aligning the ledgers

In a globally unitary universe, the most important fact about entropy is that at the finest level it does not behave the way everyday intuition says it does. If the universal state is pure, then the von Neumann entropy of the whole universe is:

$S(\rho_{universe}) = 0$

and it stays 0 forever under unitary evolution. Even if the universal state is mixed, its spectrum is preserved by unitary conjugation, so its fine-grained entropy is conserved.

So where does the "entropy increase" of thermodynamics come from? From two operations internal observers constantly perform, whether they realize it or not:

1. taking subsystems (tracing out degrees of freedom), and
2. coarse-graining (lumping microstates into macrostates).

These operations produce entropies that can increase even in globally reversible dynamics. The key is to keep the ledgers separate and then show how they relate.

Entanglement entropy (subsystem entropy)
For a subsystem A with reduced state $\rho_A(t)$, the entanglement entropy is:

$S_A(t) = -Tr[\rho_A(t) \log \rho_A(t)]$

This entropy measures how entangled A is with its complement when the total state is pure, or more generally measures mixedness of A's reduced state. $S_A(t)$ can increase as entanglement spreads, even when the global fine-grained entropy is constant. In this sense, entanglement entropy is the bookkeeping variable for "information leaving A into correlations with the rest."

Thermodynamic entropy (macrostate entropy)
Thermodynamic entropy is a statement about how many microstates correspond to the same macro-description (energy, volume, temperature, etc.). It depends on a choice of macrostates and on what degrees of freedom are ignored. In the full-state setting, we can define thermodynamic entropy as an entropy of a coarse-grained description extracted from $\rho(t)$. Different coarse-grainings yield different entropies.

To align the ledgers, we identify a standard relationship:

- Entanglement entropy tracks how information becomes nonlocal and inaccessible to small subsystems.
- Thermodynamic entropy tracks how coarse macro-descriptions fail to capture microscopic correlations and therefore treat many microstates as equivalent.

In practice, the arrow of time in thermodynamics often correlates with growth of entanglement between “system” and “environment” and with the spreading of initially simple correlations into complex multi-body correlations that are invisible to coarse variables. The universe becomes harder to describe in low-dimensional macro variables, so entropy rises.

Complete access lets us test this correspondence directly. For any proposed macro-variable set (the candidate “thermodynamic description”), we can compute:

- how much information those variables capture about the microstate,
- how quickly that captured information declines as correlations spread,
- how entanglement growth correlates with thermodynamic entropy increase.

The result is an exact accounting: thermodynamic irreversibility is not a violation of unitarity, but a loss of recoverability when restricted to certain macrostates and local subsystems.

9.2 Low-entropy boundary conditions and the growth of accessible correlations

The classical story of the arrow says: the universe began in a very special low-entropy state, and from that special beginning, typical evolution leads to increasing entropy. Without special initial conditions, the arrow would not exist.

In the full-state setting, we can state this as a testable claim about the universal record. We can ask:

- At the earliest times accessible in the dataset, does the universal state exhibit unusually low entropy relative to the natural macrostates of the later universe?
- Do correlations begin in a simple, compressible configuration and then spread into complex, high-order entanglement structures?
- Is there a temporal asymmetry in compressibility that picks out one direction as “past”?

This reframes “low entropy initial condition” as “high compressibility boundary condition.” A low-entropy macrostate is not merely “ordered”; it is describable with few parameters. A high-entropy macrostate requires many parameters or looks typical relative to a macro-description.

With complete access, we can quantify the “specialness” of boundary conditions in multiple ways:

- by coarse-grained entropy in a natural macrostate partition,
- by algorithmic compressibility of the state’s correlation structure,

- by the distribution of mutual informations among subsystems,
- by the presence or absence of stable records and redundant encodings.

A particularly illuminating quantity is **accessible correlation**: how much of the universe's correlation structure is available to modest subsystems and can be harvested as work or prediction. Early in a low-entropy universe, correlations may be structured in simple, accessible ways: temperature gradients, low entanglement across macroscopic cuts, and free energy in organized forms. As time evolves, correlations may remain present but become "locked" into fine-grained entanglement and high-order correlations that are inaccessible to typical agents. The world still contains information, but it becomes unexploitable for macroscopic control. That is irreversibility as experienced.

So the arrow becomes: growth of inaccessible correlations. The universe's information is not destroyed; it is buried in complexity.

The dataset would show this as a time-asymmetric shift:

- from low-order, macroscopic correlations that are easy to summarize and exploit,
- to high-order, delocalized correlations that require enormous global knowledge to reverse or harness.

In other words, the arrow is a monotone in "how much of the correlation structure is accessible under local coarse-grained description." With complete access, we can compute this monotone explicitly and test whether the early boundary condition is indeed special.

9.3 Coarse-graining choice as a physical question: which macrostates are natural

A common critique of entropy arguments is that entropy depends on coarse-graining, and coarse-graining depends on us. If we choose a weird macrostate partition, entropy might decrease; if we choose another, it increases. Does that make the arrow subjective?

The full-state thought experiment forces a sharper answer: coarse-graining is not arbitrary if we demand compression, stability, and record formation. The "natural" macrostates are those that:

- are stable under the dynamics (they map into one another in a simple way),
- correspond to variables that become redundantly recorded (classical facts),
- support efficient prediction of subsystem behavior,
- align with locality (macrostates defined by local densities and currents),

- and minimize description length for large-scale behavior.

In other words, macrostates are not chosen by human taste. They are selected by the universe's own dynamics and record architecture. Temperature, pressure, and density become "natural" because they are exactly the kind of coarse variables that (i) evolve predictably under local interactions and (ii) are what the environment tends to encode and propagate.

With complete access, this becomes a definable selection problem: search over families of coarse-grainings and score them by compressive adequacy. The best coarse-grainings are those that permit a low-dimensional effective dynamics with minimal loss of predictive power.

This turns the arrow-of-time debate into a concrete criterion:

- The "real" thermodynamic entropy is the entropy of the macrostate partition that is selected by maximal compression subject to robustness and locality constraints.

On that view, entropy increase is not merely a feature of any coarse-graining. It is a feature of the physically privileged coarse-grainings that yield stable classical description. If a bizarre coarse-graining makes entropy decrease, that is not a counterexample; it is an irrelevant coordinate choice that does not track any stable record structure or predictive macro-regularity.

So coarse-graining becomes a physical question because the universe itself supplies criteria for which coarse-grainings correspond to stable, objective macroscopic variables. The full dataset allows us to identify those coarse-grainings rather than assume them.

9.4 Could an arrow emerge without special initial conditions? What the data would show

The most provocative question is whether a time arrow can emerge without a special low-entropy boundary condition. Some proposals suggest that typicality arguments, internal observer selection, or dynamical attractors could generate an apparent arrow even in a time-symmetric universe. Others insist that without special initial conditions, there is no arrow—only fluctuation and symmetry.

With full access, we can answer this empirically by looking for signatures that distinguish "boundary-driven arrow" from "emergent arrow without special boundaries."

Here is what the dataset would show in each case.

Case A: Arrow driven by special low-entropy boundary
We would expect:

- a clear asymmetry in coarse-grained entropy relative to the privileged macrostate partition: low in one temporal direction, higher in the other;

- growth of entanglement and complexity away from that boundary in the “future” direction;
- a corresponding asymmetry in record formation: stable records accumulate in one direction (memory), not the other;
- compressibility asymmetry: the state is describable by few macroscopic parameters near the boundary and becomes increasingly complex later.

This is the familiar story, now expressed in measurable features of $|\Psi(t)\rangle$ and $\rho(t)$.

Case B: Arrow without special initial conditions (symmetric or “two-sided” arrow)
We would expect something different:

- either the entropy is low in a central region and grows in both temporal directions (a “Janus” or two-headed arrow),
- or the entropy is high almost everywhere with occasional low-entropy fluctuations that create localized arrows around them.

In such a world, “the arrow” would be parochial: it would exist for observers located near a low-entropy fluctuation because their memory and record formation would point away from that fluctuation in the direction of increasing entropy. But globally, there would be no unique direction singled out; there would be many local arrows.

The dataset would reveal this by showing:

- no globally privileged low-entropy past across the entire universe,
- instead, scattered pockets where compressibility is unusually high and entropy unusually low, with arrows pointing away from those pockets,
- and potentially a symmetry where for every arrow region pointing one way, there are arrow regions pointing the other way in other parts of the history.

Case C: No arrow at all (steady maximal entropy)
We would see:

- coarse-grained entropy near maximal almost everywhere,
- minimal free energy gradients,
- no stable accumulation of records in one direction more than the other,
- memory-like structures would be rare and short-lived, because record formation itself requires low-entropy resources and stable redundancy.

In such a universe, observers like us would be exceedingly unlikely, and the full record would show that plainly: stable, long-lived, redundantly encoded classical histories would be absent or extremely rare.

So could an arrow emerge without special initial conditions? The dataset could settle that by inspecting whether there is a global boundary condition asymmetry or merely local fluctuation-driven arrows.

This also clarifies a key logical point: even if the arrow is locally emergent around fluctuations, the existence of those low-entropy pockets is still a kind of special condition—just not a single global initial boundary. The deeper question is what the full record says about the distribution of such pockets and whether their existence is typical under the inferred dynamics.

In summary, the full-state thought experiment lets us align the entropy ledgers exactly: global fine-grained entropy can be constant while subsystem entanglement entropy and coarse-grained thermodynamic entropy increase. The arrow of time becomes the growth of inaccessible correlations and the irreversible accumulation of stable records under privileged coarse-grainings selected by compression and robustness. Low-entropy boundary conditions translate into high-compressibility conditions visible in the early record. And the question of whether the arrow requires special initial conditions becomes an empirical classification problem: does the dataset show a single global low-entropy boundary, a two-headed arrow, local fluctuation-driven arrows, or near-maximal entropy everywhere?

10. The Renormalization Ladder: Deriving Effective Laws at Every Scale

If the paper's central claim is that "laws are the best compression of history," then renormalization is the machinery that turns that claim into a multi-scale program. The universe does not come with a single, human-sized set of variables. It contains structure across many orders of magnitude, from subatomic degrees of freedom to planets, cells, brains, and societies. A single microscopic generator may exist, but it is not the form in which most of the world is predictable. What makes the world tractable is that *different effective laws become valid at different scales*, and those laws can be derived—at least in principle—by systematically "integrating out" irrelevant detail.

In ordinary practice, renormalization and effective field theory (EFT) are already a kind of controlled compression: you keep the degrees of freedom relevant to a scale and ignore the rest, adding effective terms that summarize what the ignored degrees do. But our ordinary practice is constrained by limited experimental access and by the impossibility of tracing all

correlations exactly. In the full-state thought experiment, those limitations vanish. We can perform renormalization not as an art but as an exact operation on ρ(t) and on the inferred generator.

This section treats renormalization as the universe's own ladder of compressions: a hierarchy of descriptions that become progressively coarser but also progressively more stable, more classical, and more useful for prediction. "Effective" becomes precise because we can compute the error introduced by coarse-graining at each step. Universality classes and fixed points become features of the actual state history, not merely of models. And bridging from physics to chemistry and biology becomes a question of legitimate derivation: can we show that higher-level regularities are stable consequences of lower-level dynamics under the relevant coarse-grainings, rather than post hoc narratives?

10.1 Integrating out degrees of freedom as a controlled compression

To "integrate out" degrees of freedom is to replace a detailed description with a coarser one that preserves what matters at a chosen scale. In quantum language, the most basic operation is taking a reduced state:

$\rho_S(t) = Tr_{E}[\rho(t)]$

where S is the subsystem of interest and E is what we discard. But renormalization is not merely tracing out an environment. It is a systematic transformation of description that defines new effective variables and effective dynamics at a coarser resolution.

In the full access setting, we can define a renormalization step as three coupled choices:

1. **A partition/coarse-graining map**: how micro-degrees are grouped into blocks or coarse variables.
2. **A projection rule**: which information is retained (the relevant operators) and which is discarded.
3. **An induced dynamics**: the effective generator that best reproduces the retained information over the relevant timescales.

What makes this "controlled" is that we can measure the error exactly. Given a candidate coarse-graining, we can compare:

- the true evolution of retained observables derived from ρ(t), versus
- the evolution predicted by the effective generator on the coarse description.

This allows a rigorous notion of approximation quality. In ordinary physics, we often justify coarse-graining by plausibility and empirical success. Here, we justify it by provable fidelity relative to the universal record.

Two additional criteria make the ladder meaningful rather than arbitrary.

Stability
A good coarse-graining yields macrostates that are stable under dynamics: they do not fragment immediately into micro-sensitive chaos; they evolve predictably in a low-dimensional way.

Closure
A good coarse-graining yields an effective description in which the retained variables are approximately closed under time evolution. If the effective dynamics constantly leaks into discarded variables in ways that feed back unpredictably, then the coarse description is not legitimate at that scale.

So "integrating out" becomes: remove degrees of freedom in a way that preserves predictability and record-structure, and quantify the residual error. That is compression with receipts.

10.2 Exact EFT emergence: why "effective" becomes precise in this experiment

In normal EFT language, we say: at low energies, high-energy details don't matter much; they renormalize couplings and generate suppressed higher-order operators. We treat this as a principle supported by experiments and by the structure of local quantum field theories. But we rarely have the ability to compute the full mapping from UV to IR exactly. "Effective" means "good enough."

In the thought experiment, "effective" can be made precise in at least three ways.

(i) Exact matching of correlation functions
Because we know ρ(t), we can compute any correlator of the retained degrees of freedom exactly. An EFT at a given scale is legitimate if it reproduces those correlators within a specified tolerance over a specified domain (time range, energy range, spatial scale). This is no longer a matter of qualitative agreement; it is an exact numerical statement about matching the universe's correlators.

(ii) Exact error bars and breakdown scales
An EFT always has a regime of validity. Usually we infer breakdown scales by dimensional analysis and by seeing where predictions fail. Here we can compute breakdown scales directly by measuring where the mismatch between EFT predictions and the true reduced dynamics exceeds a threshold. Every step of the ladder comes with a certified domain of validity.

(iii) Explicit induced operators from integration
When degrees of freedom are integrated out, new effective terms appear. In ordinary practice, we list all terms allowed by symmetry and fit coefficients. Here we can infer coefficients from the universal record, and we can also check whether some terms are genuinely absent. The operator content of the EFT becomes an empirical deduction: the minimal set of operators needed to reproduce the retained sector's dynamics.

This turns the EFT program into a compression algorithm:

- Choose a scale and a set of retained variables that correspond to what is accessible at that scale.
- Derive the minimal generator that reproduces their observed dynamics with minimal complexity.
- Record the induced couplings and the certified error.

In that sense, "effective" becomes a rigorously bounded statement: a theorem about how well a reduced model compresses the full history on a given domain.

One deeper payoff follows. Many debates about "fundamental" versus "effective" are really debates about where the compression stops. In the full-access setting, we can map the entire ladder and see whether there is a stable fixed "bottom" description that remains simplest, or whether descriptions continue to reorganize without a single privileged base. Renormalization becomes not just a computational tool but an ontology diagnostic: it tells us whether the universe has a clean fundamental layer or is better understood as a stack of mutually consistent effective layers.

10.3 Phase transitions, universality classes, and fixed points as features of $|\Psi(t)\rangle$

Phase transitions and universality classes are among the most striking examples of emergence: very different microscopic systems flow to the same macroscopic behavior. That is exactly the kind of pattern our paper treats as "law as compression." Universality says that the macroscopic law forgets microscopic detail.

In the full-state setting, universality is not merely a theoretical expectation. It is something we can detect directly in $|\Psi(t)\rangle$ by analyzing how coarse-grained descriptions change with scale and how families of systems converge under renormalization.

Fixed points as compression attractors
A renormalization fixed point is a description that looks the same under coarse-graining. In ordinary RG, fixed points correspond to scale-invariant theories and often to critical behavior.

Here, we can search for fixed points by applying the coarse-graining transformations to the inferred effective generators and checking whether the generator parameters converge.

Algorithmically:

- take many instances of subsystems across the universe that realize similar phenomena at different scales,
- coarse-grain them repeatedly,
- infer effective generators at each step,
- examine whether the parameter flows converge to stable attractors.

If they do, those attractors are fixed points—deduced not from a toy model but from the universe's own multi-scale record.

Universality classes as equivalence under coarse-graining

Two microscopic systems are in the same universality class if, after coarse-graining, their effective descriptions converge to the same fixed point behavior (same critical exponents, same scaling functions). With complete access, we can define universality classes empirically: sets of subsystems whose coarse-grained dynamics become indistinguishable within a tolerance under the renormalization ladder.

This is the ultimate form of "forgetting details." The dataset would show that many micro-differences do not matter for macroscopic predictions, because they renormalize away. That is exactly what makes science possible at human scales.

Phase transitions as reorganizations of the effective variable set

Earlier sections treated phase transitions as reorganizations in correlation and compressibility structure. In the renormalization ladder, a phase transition is where the best coarse variables change discontinuously: order parameters appear, symmetries break, correlation lengths diverge, and a new effective description becomes dramatically more compressive.

Because $|\Psi(t)\rangle$ is fully known, we can detect phase transitions not only by macroscopic observables but by state-structure signatures:

- changes in correlation-length scaling,
- changes in entanglement scaling (area laws versus critical scaling),
- abrupt changes in which coarse-grained operators are relevant.

So fixed points and universality classes become the universe's compression landmarks: places where the ladder stabilizes, and where many micro-stories collapse into one macro-law.

10.4 Bridging to chemistry and biology: what counts as a legitimate derivation

The most ambitious move is to say: if we have the full quantum record and infinite compute, we can derive not only physics at larger scales but chemistry and even biology. That claim needs discipline, because "derive biology from physics" is often used either as a triumphalist slogan or as a reductionist cudgel. The paper's stance is narrower and more rigorous: a derivation is legitimate only if it shows that higher-level regularities are stable consequences of lower-level dynamics under the relevant coarse-grainings, with quantified error and clear domain of validity.

So what counts as a legitimate derivation?

(i) Identification of stable effective degrees of freedom
For chemistry, this might be atoms, orbitals, bonds, vibrational modes, reaction coordinates. For biology, it might be macromolecules, membranes, regulatory networks, metabolic pathways. These are not arbitrary categories; they must emerge as *stable* structures in the renormalization ladder: patterns that persist, recur, and support predictive closure.

(ii) Demonstration of autonomy and closure at the higher level
A higher-level law is legitimate if it is approximately closed: knowing the higher-level state variables is enough to predict their evolution to the desired accuracy without tracking the microstate in detail. This is the hallmark of an effective theory. In chemistry, reaction rate laws and equilibrium thermodynamics have this autonomy. In biology, certain network motifs and homeostatic dynamics have partial autonomy. The derivation must show where autonomy holds and where it fails.

(iii) Quantified error and explicit breakdown conditions
A legitimate derivation cannot hide behind "in principle." It must specify: under what conditions does the higher-level description track the universal record closely, and under what conditions does it break? For example, classical chemistry breaks down in certain quantum regimes; biological models break down outside certain environmental conditions or when rare molecular events dominate.

(iv) Compression advantage as the justification for level choice
Why use chemistry rather than particle physics to predict a reaction? Because chemistry is a vastly better compression at that scale. A derivation is legitimate if it demonstrates that the higher-level variables are not merely convenient but *optimal* relative to a compression criterion: they capture the relevant predictive structure with minimal complexity.

This framework also clarifies what the full access does *not* guarantee. It does not guarantee that everything is cleanly reducible to a single neat story. Biology, especially, involves historical contingency, path dependence, and selection—features that are not merely microphysical but

organizational. Even with full access, you may find that biological regularities are less universal and more context-dependent than physical ones. The ladder might become messier: fewer fixed points, more branching, and more dependence on boundary conditions.

But "messier" is not "non-derivable." It simply means that the best compression at biological scales may involve richer state descriptions and more conditional laws. The full dataset allows that distinction to be made honestly: we can see exactly how much of biology is captured by stable effective laws and how much is contingent history encoded in the particular trajectory the universe took.

So the bridge to chemistry and biology, in this paper's framework, is not a philosophical reduction. It is a graded program of multi-scale inference:

- derive atoms as stable effective excitations,
- derive bonds and reaction pathways as constrained dynamics in reduced spaces,
- derive biochemical networks as higher-level effective dynamics with bounded error,
- and identify where the ladder yields robust laws versus where it yields only historical description.

That is what "legitimate derivation" means here: not total reduction, but certified compression across scales, with explicit autonomy and explicit limits.

In summary, the renormalization ladder is the paper's multi-scale engine: integrating out degrees of freedom becomes controlled compression with exact error accounting; "effective" becomes precise because correlators, breakdown scales, and induced operators can be computed directly from the universal record; phase transitions and universality classes become empirical features of $|\Psi(t)\rangle$ visible as RG fixed points and parameter flows; and bridging to chemistry and biology becomes legitimate when higher-level variables and laws are shown to be stable, autonomous compressions of the lower-level dynamics with quantified domains of validity.

11. Counterfactuals and Contingency: What "Would Have Happened" Can Mean

Complete access to the universal state history tempts an intoxicating conclusion: if we know everything that did happen, we can compute everything that would have happened. That temptation is understandable—and partly justified—but it needs discipline. A realized history is one trajectory. Counterfactuals require a space of alternatives. To speak responsibly about

"would have happened," we must specify what is held fixed, what is allowed to vary, and what rule generates the alternatives.

This section builds a rigorous notion of counterfactuals appropriate to the paper's framework. Counterfactuals are not mystical; they are the lawful consequences of a reconstructed generator under perturbed conditions. But the counterfactual leap is not automatic: from one history, the generator may be underdetermined, and factorization choices may be non-unique. Moreover, in a quantum setting, "perturbing the state" is not as simple as changing a classical variable; different perturbations may not correspond to realizable operations within the universe.

So the paper's posture is: counterfactuals are *model-dependent but testable within the model*, and their legitimacy depends on the stability of the inferred generator and on robustness under reasonable representation choices. Within that discipline, complete access unlocks three major capabilities: sensitivity analysis, identification of control knobs and bottlenecks, and explicit mapping of the limits—where counterfactual talk dissolves into coordinate choice.

11.1 From one realized history to a lawful ensemble: the counterfactual leap

A counterfactual is a statement of the form:

- "If condition C had been different at time t, outcome O would (or would likely) have followed."

To make that statement meaningful, we need two ingredients:

1. a rule of evolution (the generator or dynamical law), and
2. a definition of the alternative condition C' (the perturbed initial/boundary condition).

Given those, a counterfactual is simply the evolved consequence of C' under the same law.

But here is the key: the universal record gives us **one** realized path. It does not directly provide the distribution over alternative initial conditions, nor does it guarantee that the inferred generator is unique. Therefore, the counterfactual leap is a methodological step:

- we infer the simplest generator consistent with the record (subject to locality, symmetry, etc.), and
- we define a family of nearby states or boundary conditions that count as "plausible alternatives," and then
- we propagate them forward to build a lawful ensemble of alternative histories.

This is exactly where the paper's compression principle becomes decisive. The lawful ensemble is not arbitrary; it is the set of trajectories generated by the most compressive inferred law when applied to a neighborhood of perturbations. In that sense, "would have happened" means "would have followed under the minimal generator that compresses the actual history."

However, there is a subtle but important distinction between two kinds of counterfactual ensembles:

(i) Perturbation ensembles (nearby histories)
We vary the state slightly at some time slice or boundary and see what changes. These counterfactuals answer questions like: "if a single molecule had been displaced, would the hurricane still form?" or "if a neuron had fired differently, would the decision change?"

(ii) Re-sampling ensembles (alternative microstates consistent with the same macrostate)
We hold fixed certain macroscopic conditions and vary the microscopic realization within that macrostate. These counterfactuals answer: "given the same macrostate, what range of micro-histories are typical?" This connects counterfactuals to statistical mechanics and to the arrow of time: many microstates correspond to one macrostate, and typical evolution produces typical macroscopic outcomes.

Complete access helps because it lets us define both ensembles explicitly. We can specify exactly what macro constraints are being held fixed and then sample microstates consistent with them, rather than relying on heuristic distributions.

But the leap remains: we are moving from one history to many possible histories. That move is justified only insofar as we have a stable generator and a principled perturbation rule.

11.2 Sensitivity analysis: micro-perturbations and macro-outcomes

Once we have a generator and a perturbation ensemble, we can do sensitivity analysis: quantify how much macroscopic outcomes depend on microscopic details.

In classical chaos, the lesson is that micro-perturbations can lead to divergent trajectories in phase space, even while macroscopic statistical regularities remain stable. In quantum systems, the story is more nuanced: unitary evolution preserves inner products globally, but local subsystems can become effectively unpredictable because information disperses into entanglement and becomes inaccessible. Sensitivity can therefore show up as:

- rapid growth of distinguishability between reduced states of macroscopic variables under micro-perturbations,
- changes in which branch records become realized in an objective-collapse scenario,

- or stability of coarse-grained statistics despite micro-level divergence.

With full access, sensitivity analysis becomes exact and multi-layered.

Define macro-outcomes precisely
First we must define what "macro-outcome" means in this framework: it is a stable record variable in some subsystem (Section 8) or a macrostate in a privileged coarse-graining (Section 9–10). Examples include: "this pointer reads 1," "this region's temperature rises," "this organism survives," "this social structure persists." The definition is not rhetorical; it is a function of the state's record structure.

Define perturbations precisely
Then define a family of micro-perturbations at time t: small changes in a localized region's state while keeping the rest fixed in a controlled way. In quantum terms, this might mean applying a small unitary supported on a local region, or perturbing the reduced density matrix in a way that is consistent with a global state. In the thought experiment, we can construct these perturbations at will within the simulation.

Compute response functions and divergence rates
Now we propagate each perturbed state forward and compute how macro-outcome variables differ. We can define sensitivity metrics such as:

- the probability (or branch-weight measure) of each macro-outcome as a function of perturbation amplitude,
- the time-to-divergence for macro-variables (how long two perturbed histories remain macro-indistinguishable),
- and the stability of statistical outcomes across ensembles.

This produces a quantitative answer to "contingency." Some macro-outcomes may be extremely robust: nearly all micro-perturbations yield the same macrostate. Others may be knife-edge: tiny perturbations flip the outcome with substantial frequency.

A key conceptual payoff emerges: the universe can be micro-chaotic yet macro-stable. Sensitivity analysis reveals which layers of description are legitimate. If macro-outcomes are stable under micro-perturbations, that validates the effective laws at that scale (Section 10). If macro-outcomes are not stable, then that scale lacks autonomous lawfulness and is dominated by historical contingency.

So sensitivity analysis is the bridge between "deterministic micro-law" and "contingent macro-history." It identifies where history is compressible into laws and where it is not.

11.3 Control knobs and bottlenecks: rare degrees of freedom with outsized leverage

One of the most practical things an omniscient analysis could deliver is the identification of leverage points: degrees of freedom that are rare or localized but have disproportionate influence on macro-outcomes. In complex systems language, these are control knobs, bottlenecks, and tipping points.

In the paper's framework, a "control knob" is defined by counterfactual leverage:

- A localized variable K is a control knob for outcome O if small perturbations of K produce large changes in the distribution of O, relative to perturbations of other comparable-sized variables.

A "bottleneck" is a structural choke point:

- A variable or subsystem B is a bottleneck if many causal pathways and influence flows must pass through it to affect downstream macro-outcomes.

Complete access enables a principled search for these.

Leverage via influence gradients
Using the influence graph (Section 6), we can compute how changes propagate. Some nodes may have high out-degree in the influence graph or occupy central positions in many shortest influence paths. These are candidates for bottlenecks.

Leverage via sensitivity ranking
Using sensitivity analysis (11.2), we can rank local perturbations by their effect on macro-outcome distributions. Most perturbations will wash out; a few will matter. Those few define control knobs.

Leverage via renormalization relevance
There is a deep connection to renormalization: relevant operators are exactly those perturbations that grow under coarse-graining and dominate large-scale behavior. In this sense, control knobs are the physical instantiation of "relevant directions" in the renormalization flow. A perturbation with outsized leverage is one that does not renormalize away.

This yields a unifying picture:

- Bottlenecks are where influence concentrates.
- Control knobs are perturbations aligned with relevant directions.
- Contingency often arises at points where the system sits near a critical boundary between macroscopic regimes (tipping points).

This can also explain why some historical events feel decisive. In a system near a bifurcation, small differences can determine which basin of attraction the macro-dynamics enters. Full-state counterfactual analysis would identify such bifurcation points precisely: times and regions where macro-outcome distributions are highly sensitive to particular degrees of freedom.

But a caution matters. "Leverage" is always defined relative to a choice of macro-outcome and a choice of coarse-graining. A molecule can be a control knob for one outcome and irrelevant for another. So the right claim is conditional: given a specified macro-description, the analysis identifies the degrees of freedom with outsized causal leverage for that description.

11.4 Limits: counterfactuals under dualities and non-unique factorizations

The paper must explicitly state where counterfactual talk stops being well-defined, even with infinite compute and complete data. Two limits are especially important.

Limit 1: Dualities and representational equivalences
If two theoretical descriptions are related by a duality—different variables, same invariant content—then counterfactuals phrased in duality-dependent terms may be meaningless. For example, "what if field X had been larger" may not be an invariant question if X is not a uniquely defined physical quantity across equivalent representations. The only counterfactuals that survive dualities are those formulated in terms of invariant structures: observable correlators, conserved charges, record variables, and other representation-independent features.

So the rule is:

- Legitimate counterfactual claims should be stated in terms of invariants of the equivalence class of generators compatible with the history.

If a counterfactual depends on a gauge choice, it is not a fact about the world; it is a fact about a coordinate system.

Limit 2: Non-unique factorization of the universal Hilbert space
Many of our counterfactual notions rely on "local perturbations" of "subsystems." But subsystem decomposition is not uniquely given by the mathematics. If factorization is ambiguous, then the notion of "local change" can vary with representation. A perturbation that is local in one factorization can be highly nonlocal in another.

This threatens causal and counterfactual inference unless we constrain factorization choices. The paper's consistent solution is to privilege factorizations that maximize locality and compressibility of the generator and that align with stable record structures. Within that

privileged class, counterfactuals are meaningful. Outside it, counterfactuals may be coordinate artifacts.

Therefore, the limit is:

- Counterfactuals are well-defined only relative to a representation class selected by compression, locality, and record robustness. If multiple inequivalent representations compress equally well, some counterfactual questions may be underdetermined.

Limit 3: Generator underdetermination from a single trajectory
Even with full access, the inferred generator may not be unique (Section 3.4). Different generators can reproduce the realized history but differ on counterfactual trajectories. In that case, "what would have happened" is not a single-valued answer; it is model-dependent. The correct output is then a family of possible counterfactuals, with the invariant intersection treated as deduced and the remainder labeled as underdetermined.

This is an important honesty clause: omniscience about the realized past does not automatically grant omniscience about all possible futures, unless the law is uniquely determined.

In summary, counterfactuals in this experiment become lawful simulations under an inferred minimal generator applied to principled perturbation ensembles. Sensitivity analysis quantifies where macro-outcomes are robust versus contingent under micro-variation. Control knobs and bottlenecks are rare degrees of freedom aligned with relevant directions in the renormalization flow and central in the influence graph. And the limits are real: counterfactual claims must be stated in invariant terms to survive dualities, must be anchored in representation classes selected by locality/compression to survive factorization ambiguity, and may remain underdetermined if the generator is not uniquely fixed by the realized history.

12. Underdetermination and Duality: When Perfect Data Doesn't Pick One Story

The thought experiment is designed to dissolve many philosophical disputes by replacing them with computations. But not all disputes dissolve. Even with perfect, total data—$t \mapsto |\Psi(t)\rangle$ or $t \mapsto \rho(t)$—there can remain genuine underdetermination: multiple, inequivalent *stories* that compress the record equally well and agree on every operationally accessible prediction. This is not a failure of analysis. It is a feature of how description works.

Physics has already taught us this lesson in miniature. Gauge theories admit multiple mathematically distinct descriptions of the same physical situation. Dualities relate theories that look different but yield the same observable content. Even in classical mechanics, different

coordinate systems can make the same motion look radically different. In our "universal dataset" setting, the stakes are higher because the temptation is to treat the dataset as an oracle that selects a single ontology. But if two descriptions are truly equivalent—if they generate the same invariants and the same record-structure—then the dataset cannot privilege one as "what the world really is," except by extra criteria that go beyond empirical content.

This section articulates that point carefully. Dual theories can be compression-equivalent. Subsystem factorization can be ambiguous. Ontology can float while invariants remain fixed. The paper's response is not despair but discipline: the proper output of the analysis is a set of invariants and operational reconstructions that are stable across equivalent descriptions, plus a clear statement of where alternative ontologies remain permissible.

12.1 Dual theories, identical predictions: the compression-equivalence problem

A duality, in the relevant sense, is an equivalence between two theoretical descriptions that map states to states and observables to observables in a way that preserves all measurable content. The two theories can use different fundamental objects—fields versus strings, particles versus vortices, geometry versus entanglement networks—yet still agree on every observable correlation and dynamical prediction.

In the full-state thought experiment, dualities show up as a subtle type of underdetermination: two different inferred generators, expressed in two different variable sets, can reproduce exactly the same universal history. If both generators yield the same compressibility score under the paper's criterion—minimal description length subject to locality/symmetry/robustness constraints—then the data alone does not choose between them.

This is the **compression-equivalence problem**:

- If two descriptions compress the full record equally well and are linked by an exact equivalence mapping, then "the best compression" is not unique. The world permits more than one minimal story.

This is not a mere philosophical quibble. It can appear concretely in at least three ways.

(i) Exact duality mappings
There exists a bijection between the descriptions such that every prediction and every invariant is preserved. Then the dataset cannot break the tie, because breaking it would require privileging one set of primitives over another without empirical leverage.

(ii) Approximate dualities in restricted regimes
Two descriptions might be indistinguishable across the domain sampled by the realized history

but differ in counterfactual regimes not explored. In that case the dataset may still underdetermine the true law, and counterfactual claims become model-dependent. The underdetermination is not fundamental; it is contingent on what the universe happened to do.

(iii) Multi-level equivalences
A microscopic description and an effective macroscopic description can both compress the data well but in different ways: one is fine-grained and uniform across all scales; the other is coarser but simpler in the regimes that dominate the history. If both describe the record within tolerance, choosing between them becomes a pragmatic decision about which questions you want to answer, not a decision about what the world really is.

In all these cases, perfect data does not force a single metaphysical conclusion. It forces a set of equivalence classes of descriptions. The correct scientific stance is to report the class and the invariants shared by the class, rather than pretend the class has a unique representative.

12.2 Subsystem factorization ambiguity: "what is a system?" as a choice

Much of the paper's analysis—locality, causality, decoherence, classicality, renormalization—relies on talking about subsystems: system A, environment E, fragments F, and so on. But the universal Hilbert space does not come with a unique, God-given tensor product decomposition into "things." Factorizations can be chosen in different ways, and different choices can make different structures look local or nonlocal, classical or quantum, simple or complex.

This creates a second kind of underdetermination: not about the law, but about the ontology of parts.

The question "what is a system?" becomes nontrivial in this setting. If you can factor the same global state space into different subsystem structures, then statements like:

- "A interacts locally with B,"
- "A decoheres due to E,"
- "A contains a record of B,"

depend on the chosen factorization. With complete access, we can compute everything relative to any factorization, but we still face the question of which factorization is physically meaningful.

The paper's answer is not "any factorization is as good as any other." In practice, the world strongly constrains useful decompositions, because only some decompositions yield:

- a local generator with sparse coupling structure,

- stable classical records and pointer variables,
- robust macroscopic variables with predictive closure,
- a coherent renormalization ladder.

But ambiguity can remain if multiple inequivalent decompositions satisfy these constraints equally well.

There are at least three relevant categories:

(i) Trivial re-labelings
Factorizations related by simple permutations or local basis changes that do not alter what counts as local structure. These are harmless and can be treated as gauge-like.

(ii) Dual factorizations
Factorizations that reorganize degrees of freedom nontrivially but still preserve all operationally accessible structures (locality, records, etc.) under an equivalence. These are more serious: the "parts" of the world are not uniquely fixed.

(iii) Pathological factorizations
Factorizations that scramble locality and record structure. These are mathematically allowed but physically unhelpful because they destroy compressibility and do not align with any stable observational architecture.

Perfect data allows us to identify which category we are in. But it does not guarantee uniqueness. If the universe's structure supports more than one "natural" factorization, then the ontology of subsystems is underdetermined by the record.

This is the most direct way perfect data fails to pick a single story: even if we know everything, the carving of the world into "things" may remain a representational choice within an equivalence class.

12.3 Ontology versus invariants: what remains fixed across equivalent descriptions

If descriptions can differ while preserving the same empirical content, what remains fixed? The answer is: invariants. Invariants are the elements of description that survive transformation between equivalent theories or equivalent factorizations.

In the context of this paper, the most important invariants are:

(i) Correlational invariants
Full sets of correlation functions among operationally defined observables (or their

equivalents), and the relational patterns that define record structures. If two descriptions produce exactly the same correlator structure for all operationally meaningful queries, then they are empirically equivalent.

(ii) Dynamical invariants
The equivalence class of time-evolution maps as they act on the invariant observable algebra. Even if one description uses Hamiltonians and another uses different primitives, what matters is the induced evolution on the invariants.

(iii) Symmetry and conservation structure
Global symmetry groups and conserved charges that can be reconstructed operationally. While the representation of a symmetry may differ across duals, the invariant content—selection rules, conserved sectors, and transformation properties of records—remains stable.

(iv) Causal/influence constraints
Operational speed limits (no-signaling, finite propagation of influence in the privileged reconstruction), causal ordering relations among records, and horizon-like constraints. These are the kinds of structures that survive re-description if they are truly physical.

(v) Compression minima and equivalence classes
Even when there are multiple minimal stories, the minimal description length itself and the set of all descriptions that achieve it are invariants of the compression problem. The universe may not choose one story, but it does constrain the space of equally good stories.

Ontology—what the world is "made of"—can float while invariants remain fixed. This is not a defeat. It is a clarification. The stable content of science is the invariant structure; ontology is the set of human-friendly narratives that realize that structure.

This has an important practical implication for the thought experiment. If we ask "what could we deduce?" the strongest deductions will be invariants. If we ask "what is really real?" we may not get a unique answer. We may get an equivalence class of answers that all compress the record equally well.

12.4 A practical rule: prefer invariants and operational reconstructions

Given underdetermination, the paper proposes a practical discipline for how to report conclusions from the universal dataset:

- Prefer invariants over ontological claims.
- Prefer operational reconstructions over metaphysical commitments.

- When multiple stories are compression-equivalent, report the equivalence class and the shared structure, and treat the choice of narrative as pragmatic.

Concretely, this rule has four parts.

(i) State results in invariant terms whenever possible
Instead of "the world is fundamentally fields," say "the invariant observable algebra evolves according to X and exhibits locality structure Y." Instead of "spacetime is fundamental," say "the influence graph admits an emergent metric embedding with curvature features Z, and this embedding is stable under the preferred equivalence class." These statements survive dualities and factorization ambiguity.

(ii) Use reconstruction procedures as definitions, not assumptions
Define locality by recovered sparse generator structure. Define classicality by stable redundant records. Define time's arrow by monotone growth of inaccessible correlations under privileged coarse-grainings. These are operational reconstructions from $|\Psi(t)\rangle$, not imported metaphysical primitives.

(iii) Track ambiguity explicitly
Where the data does not select a unique story—dual descriptions, multiple natural factorizations—say so. Provide the set of candidate minimal descriptions and specify which claims are common to all of them and which are story-dependent.

(iv) Choose narratives by utility, not by pretending they are forced
If two dual descriptions are equally valid, choose the one that answers the questions you care about most directly: the one that makes locality manifest, the one that aligns with chemistry, the one that simplifies cosmology, etc. That choice is legitimate as long as it is acknowledged as a choice within an equivalence class.

This discipline is not merely philosophical hygiene. It protects the project from overreach. The full dataset is powerful, but it does not repeal the logic of equivalence. When nature permits multiple perfect compressions, science's job is to identify the invariant skeleton shared by them.

In summary, perfect data can still leave genuine underdetermination: dual theories can be compression-equivalent, and subsystem factorization can be ambiguous. Ontological stories can therefore remain non-unique even when every operational prediction is fixed. What remains stable are invariants—correlation structure, induced dynamics on invariant observables, symmetry/conservation patterns, causal constraints, and the equivalence class of minimal compressions. The practical rule is to prefer invariants and operational reconstructions and to report ambiguity honestly when the dataset does not pick one story.

13. Computability Boundaries: Even Infinite Compute Doesn't Prove Minimality

The paper has leaned heavily on a central idea: **laws are the best compression of history**. This is attractive because it replaces metaphysical talk with something operational: given the total record, the "law" is the shortest program that reproduces it (or reproduces it within specified tolerances at specified scales). If we could actually find that shortest program, we would have a principled notion of explanation: the law is the minimal description; everything else is commentary.

But there is a trap hidden inside this appeal. Even if we grant "infinite compute" in the thought experiment, we have not escaped the logical limits of computation. Some questions are not hard because we lack cycles; they are hard because they are **undecidable**. In particular, "Is this program the shortest possible one that generates this data?" is not a question that a general algorithm can always answer, because it collapses into the problem of proving that no shorter program exists—a problem that is deeply intertwined with the halting problem.

So this section draws a boundary around what the compression program can legitimately claim. Compression is an extraordinarily powerful guide: it can reveal symmetries, locality, effective laws, phase structure, and causal organization. But it cannot serve as an oracle that pronounces a unique, absolutely minimal law in an unrestricted sense. What it can do instead is far more honest and still extremely strong: it can establish dominance relations, stability across perturbations, and minimality within specified model classes and operational constraints.

13.1 Laws as shortest programs: the appeal and the trap

The appeal is straightforward. Suppose the universe's full record is D (the complete $|\Psi(t)\rangle$ history, or an equivalent specification). A candidate "law" is a program P that outputs D (or outputs a generator that reproduces D). Then we want the P* with minimal length:

$P^* = \text{argmin}_{P : P \to D} |P|$

This is seductively clean because it makes explanation identical to compression. The shortest program captures the deep regularities; the longer programs contain more ad hoc detail. Under this view:

- Symmetries are compressions (reuse a rule instead of listing cases).
- Locality is compression (short-range interactions instead of a dense interaction table).
- Conservation laws are compression (rule out huge spaces of possibilities).
- Effective theories are compression (keep the variables that matter at scale).

The trap is that “shortest program” is not just difficult to find; it is not, in general, something you can certify. You might find a very short program that reproduces D, but proving that no even shorter program exists is a different kind of problem. It is a universal negative across an infinite space of possible programs.

In practice, scientists do not need absolute minimality to do physics. We need laws that are simple, stable, predictive, and robust. The trap arises when we confuse these pragmatic virtues with the claim that we have identified the unique minimal generator of reality.

The thought experiment adds another subtle trap: if we *really* had the full record, we could always build a trivial program: “print the record.” That is maximal overfitting—perfect accuracy with no explanation. Compression is supposed to reject that. But then we need a principled way to measure and certify that a found program is “as compressed as possible.” That is where computability limits bite.

13.2 Why absolute Kolmogorov minimality is not generally decidable

The underlying issue is Kolmogorov complexity: the length of the shortest program that outputs a given string (or dataset) on a fixed universal Turing machine. Kolmogorov complexity is a powerful idealization because it captures what we mean by “intrinsic compressibility.” But it comes with a brutal fact:

- There is no general algorithm that, given an arbitrary dataset, computes its exact Kolmogorov complexity.

More importantly for our purposes:

- There is no general algorithm that, given a candidate program P and dataset D, can always decide whether P is the shortest program that outputs D.

The intuition for why is simple enough to state without technical flourishes. If you could decide minimality in general, you could solve the halting problem. Sketch of the linkage:

- Questions about whether a shorter description exists can be reduced to questions about whether certain programs halt (or halt within a bound), because a program that halts can output something, and non-halting behavior can be encoded into whether a compressed description exists.
- Proving that no shorter program exists requires ruling out vast classes of programs, many of which might only differ in whether they halt after long computation.

- The inability to decide halting in general infects the inability to certify minimality in general.

So “infinite compute” does not grant omniscience about minimality. Infinite compute means you can run every program for as long as you like, but it does not give you a terminating method for deciding that a program will never halt, or that no shorter program will ever produce D. The barrier is not time; it is logical.

This has two direct consequences for the paper’s compression thesis.

1. **You can search for shorter and shorter candidate laws, but you cannot, in general, prove you have hit the absolute floor.**
2. **Even if the universe’s law is short, the proof that it is shortest may be unprovable within any given formal system.**

That second point matters because “deduction” is always deduction within a formal framework. The universe may have a minimal generator, but formal mathematics may not be able to certify that it is minimal in the absolute Kolmogorov sense.

13.3 What can be proven instead: dominance, stability, and model class minimality

If absolute minimality is not generally decidable, what can we still claim in a disciplined way? Quite a lot—enough to keep the compression program scientifically potent.

(i) Dominance: comparative compression
We can often show that one model dominates another: it is shorter and fits at least as well, or it fits better at comparable length. This is how real scientific progress works: we replace longer, less explanatory descriptions with shorter, more predictive ones.

With full access, dominance comparisons can be exceptionally sharp because goodness-of-fit is not measured on a small dataset; it is measured on the entire history. So while we may not certify global minimality, we can certify that huge regions of “story space” are inferior to a given candidate law.

(ii) Stability: robustness under perturbations and re-descriptions
A powerful proxy for “lawlikeness” is stability. If a candidate generator yields correct behavior not only for the exact realized history but also for neighborhoods of perturbations (counterfactual ensembles), then it is capturing a genuine regularity rather than memorizing the record.

Similarly, if the same structure is recovered under different coarse-grainings and different reasonable subsystem decompositions, that is evidence of invariance and therefore of deep compressibility.

Stability does what absolute minimality cannot: it separates “short but brittle tricks” from “short and structurally real” descriptions.

(iii) Minimality within a restricted model class
While Kolmogorov minimality over *all possible programs* is undecidable, minimality within a constrained class often is decidable or at least tractable. For example:

- minimal local Hamiltonian given a fixed operator basis and locality radius,
- minimal Markovian quantum channel generator given a fixed noise model,
- minimal effective field theory up to some operator dimension,
- minimal causal graph given a specified intervention/counterfactual framework.

In each case, we restrict the hypothesis space to something physically motivated and operationally meaningful. Then we can ask for the best compression *within that class* and certify it with strong methods (and in this thought experiment, with complete data).

This is exactly what working physics already does, but here it becomes explicit: the claim is not “we found the absolute shortest program for the universe.” The claim is “within the physically motivated class of local generators consistent with the record, this one is minimal (or near-minimal) and stable.”

(iv) Lower bounds and incompressibility certificates for specific features
Even if we cannot compute exact Kolmogorov complexity, we can sometimes prove lower bounds for certain aspects of the data. For example:

- if the data contains a truly random component (in the algorithmic sense), no short program can compress it beyond a point;
- if certain correlations vary in ways that cannot be captured by a restricted operator set, any model in that set must grow in complexity;
- if a description requires encoding a large lookup table of contingent facts, there is an irreducible description cost.

So we can often certify that some kinds of “short laws” are impossible in specified classes, even if we cannot certify absolute minimality across all programs.

These forms of provability—dominance, stability, model-class minimality, and lower bounds—are enough to support the paper's program without turning it into a metaphysical absolutism.

13.4 A disciplined epistemology: compression as guide, not as oracle

This section culminates in an epistemological rule that keeps the project honest:

- Use compression as a guide to structure, not as an oracle of ultimate ontology.

In practice, that means four disciplines.

(i) Replace "the law is the shortest program" with "the law is the most stable, dominant compression within a physically motivated class."
This keeps the core insight (laws compress) while avoiding the undecidable claim (absolute minimality).

(ii) Separate invariants from narratives
As Section 12 emphasized, multiple stories can be compression-equivalent. Compression can reveal invariants—symmetry structure, causal constraints, conserved quantities, record architectures—without uniquely selecting a metaphysical story about what the world "is." So we report invariants first and treat ontology as optional.

(iii) Treat residual complexity as real
Some parts of the universe's history may be irreducibly contingent: they do not compress into law, except as "this particular initial condition happened." Compression teaches humility: it tells us where laws end and where history begins. The existence of irreducible contingency is not a failure; it is a fact about the data.

(iv) Keep the falsifiability posture
Even in this omniscient setting, we should maintain the habit of specifying what would count as failure for a proposed compression: instability under perturbations, inability to reproduce correlators, lack of robustness under coarse-graining, or the existence of equally short alternative generators that disagree on counterfactuals. This is the scientific posture applied to compression itself.

So the "infinite compute" thought experiment has a surprising moral: even with everything, you still don't get a divine guarantee that your shortest found law is absolutely minimal. What you do get—if you stay disciplined—is something arguably closer to real science: a hierarchy of increasingly powerful compressions, stabilized by invariants, validated by robustness, and honest about where underdetermination and undecidability prevent final metaphysical closure.

In summary, viewing laws as shortest programs is compelling but dangerous if treated as an oracle. Absolute Kolmogorov minimality is not generally decidable, so perfect data and unlimited runtime do not automatically yield a certified minimal law. What can be established instead are dominance relations among models, stability under perturbations and re-descriptions, and minimality within physically motivated model classes, along with meaningful lower bounds on compressibility for specific features. The disciplined epistemology is therefore: compression is a guide to structure and invariants, not a guarantee of unique, ultimate ontology.

14. The Residual "Why": Where Explanation Stops Being Deduction

If we really could tag every particle, record the universal state for all time, and reconstruct the best compressive laws at every scale, we would have achieved something close to maximal scientific description. We would know "what happens" with unmatched precision and we would have extraordinary control over "what would happen" under specified lawful perturbations. We would have derived locality structure, effective field theories, classical record formation, entropy growth, and the renormalization ladder, all as features of the universal record.

And yet, even at that limit, there is a residue: the question "Why this rather than that?" does not vanish. It changes shape, it becomes more carefully posed, and some versions of it become answerable. But a core "why" remains beyond deduction, not because we lack compute, but because explanation has a boundary: at some point we hit boundary conditions, equivalence classes of laws, and the difference between describing structure and justifying it.

This section is not a retreat. It is the paper's honest finishing move: to mark the point where the compression program delivers maximal explanatory power *as science*, and where the remaining questions are either meta-scientific (about why the law has the form it does) or genuinely philosophical/theological (about meaning and value). The dataset can constrain these questions, sometimes severely. But it cannot dissolve them into computation.

14.1 Why these laws rather than others: boundary conditions and meta-laws

Even if we infer an extremely compact generator that reproduces the entire history, two "why" questions remain distinct:

1. **Why this history (initial/boundary conditions) given the law?**
2. **Why this law (this generator class) rather than another?**

The first is often addressable in the language of boundary conditions. In practice, much of what we call "contingency" is simply the universe's initial condition: a particular microstate (or macrostate) at some boundary time. Even with full access, some of that may remain irreducible: the shortest description of the initial condition may still be large. We can compress it, but not necessarily to a small set of principles. In that case, the answer to "why this history?" is simply "because that is the boundary condition," plus whatever partial compression is possible.

The second question—"why this law?"—can be pushed back, but not eliminated. Suppose we discover that the best compression of the record is a local quantum field theory of a certain form. We can then ask: why a local QFT at all? Why these symmetries, these coupling constants, these particle families, these dimensions? Any answer that stays within physics typically introduces a **meta-law**: a higher-level rule that selects the law we observe. Examples include:

- a theory space with a dynamical selection principle,
- a landscape with a measure,
- a cosmological mechanism that drives parameters toward fixed points,
- or a deeper unifying theory from which the inferred law is an effective limit.

In the full-state setting, we can detect whether such meta-law structure is *required* for compression. If the best description of the record includes a mechanism that *explains* why couplings take their values (for instance, an attractor or fixed-point selection), then the "why these laws" question partially converts into "because the meta-dynamics drives them there." That would be a genuine scientific gain.

But if the most compressive law still contains arbitrary constants or contingent symmetry choices not derivable from within the theory, then "why these laws" remains a boundary-like question at the level of law itself. The law is then the compression floor, not because it is metaphysically ultimate, but because no further compression is available without adding assumptions not fixed by the record.

So the dataset can tell us whether the law's parameters are:

- **derivable** (forced by deeper structure),
- **selected** (driven by dynamics or measure effects), or
- **residual** (left as brute facts in the best compression).

Even in the first two cases, a final "why this meta-law?" may persist. Explanation can recurse, but it does not guarantee closure.

14.2 Fine-tuning, selection, and anthropic structure: what the dataset can constrain

Fine-tuning claims are often made sloppily: “constants look tuned for life,” “the universe is improbable,” “this must be designed,” “this must be a multiverse.” The full-state thought experiment forces these claims to become precise, because it gives us a way to define what “improbable” means only relative to an ensemble and a measure, and it gives us a way to test what kinds of selection effects are consistent with the record.

The dataset can constrain fine-tuning arguments in several concrete ways.

(i) Sensitivity of macroscopic structures to parameter variation
Using the counterfactual machinery (Section 11), we can vary parameters of the inferred generator—within whatever notion of lawful alternative is defined—and simulate alternative universes. We can ask: how robust are stars, chemistry, stable long-lived complexity, and self-replicating systems to small changes in constants? This yields an actual sensitivity map rather than armchair speculation.

(ii) Existence of attractors and universality that reduce tuning
Renormalization and fixed-point structure (Section 10) can explain why many macroscopic features are insensitive to microscopic detail. If the record reveals strong universality—large basins of parameter space that yield similar macroscopic complexity—then “fine-tuning” weakens. Conversely, if complex structures exist only in knife-edge regions, fine-tuning strengthens.

(iii) Evidence for multiple domains or branching parameter regions
If the record contains signatures of domain variation—regions with different effective parameters, or cosmological processes that sample different vacua—then “selection” becomes a live scientific explanation. In such a case, anthropic reasoning can become constrained rather than free-floating, because the measure over domains and the mechanism of variation can be inferred from the record itself.

(iv) Observer selection effects as record-structure constraints
Even without a multiverse, the dataset can show how “observers” (stable record-making systems) arise only under certain thermodynamic and informational conditions: low-entropy resources, stable chemistry, and robust classicality. This does not prove anthropic arguments, but it clarifies the conditional: certain kinds of universes cannot contain record-keeping observers, and therefore cannot host this kind of dataset reconstruction from within.

However, there is also a hard limit. Fine-tuning arguments require a **measure** over possible laws or parameters. The dataset can constrain a measure if it contains evidence of a generating process (a meta-law that samples parameters). But if there is no such evidence—if the best compression contains fixed constants without a sampling story—then any probabilistic fine-

tuning claim exceeds what the dataset alone can justify. You can still ask “why these values,” but you cannot honestly assign “improbability” without importing a measure.

So the dataset can sharply distinguish:

- sensitivity facts (what changes if constants change),
  from
- probability claims (how likely a given constant set is),

and it can show when the latter have a scientific footing and when they do not.

14.3 Meaning, value, and teleology: inference limits and legitimate questions

A complete physical description—even one as maximal as our thought experiment—does not automatically answer questions of meaning and value. It can, however, constrain them in important ways.

First, it can constrain **instrumental claims**: if someone says “this purpose is written into nature as a detectable tendency,” the dataset can check whether there are lawlike attractors toward certain structures, whether complexity growth is generic, whether life-like organization is stable, and whether certain kinds of order repeatedly arise. If a teleological claim is framed as a *predictable structural tendency*, it becomes empirically assessable.

Second, it can constrain **anthropological claims** about humans and minds as physical systems: the dataset can show what kinds of information-processing architectures are possible, what kinds of stability constraints exist, and how memory and agency emerge as classical record structures in biological matter.

But the dataset cannot do two things that people often want it to do.

It cannot derive “ought” from “is.”
No amount of state-tracking yields moral obligation as a deductive consequence unless moral premises are already included. Even if the dataset reveals that certain behaviors promote stability or flourishing, the step from “promotes” to “ought” is not purely deductive. It is evaluative.

It cannot certify ultimate purpose without additional axioms.
A goal-directed description can often be constructed for complex systems (agents optimize, evolution selects, organisms regulate). These are real patterns, but they are immanent teleologies—internal to systems, not necessarily cosmic intent. To infer ultimate teleology (purpose of the universe as such) from the dataset would require either:

- explicit features in the law that encode goals, or
- a justified interpretive framework that treats certain patterns as evidence of intent.

The dataset might provide raw material, but it does not force the conclusion. The same record can often be narrated as "purposeful" or as "self-organizing under selection," and Section 12's underdetermination arguments return here: multiple narratives can fit the same invariants.

So the disciplined stance is:

- Meaning and value remain legitimate questions, but their answers are not compelled by maximal physical deduction alone.
- The dataset can constrain teleology-as-structure, but not teleology-as-ultimate-intent, unless that intent is encoded in invariant, operationally reconstructible features.

This is not dismissive. It is simply an insistence that "explanation" has categories: causal/mechanistic explanation, compressive explanation, and justificatory/teleological explanation are not identical.

14.4 A modest conclusion: maximal description still leaves metaphysics intact

The final conclusion of the paper is deliberately modest, because it wants to be maximally strong where science is strong and maximally honest where science is not.

What maximal description gives us:

- a reconstructed generator (or equivalence class of generators) that compresses the universal history,
- a derived causal/influence structure, emergent geometry criteria, and a renormalization ladder,
- an exact account of decoherence, classical record formation, entropy growth, and the arrow of time,
- certified effective laws at every scale with quantified domains of validity,
- and disciplined counterfactuals within the inferred model class.

This is enormous. It is, in a sense, the completion of the scientific project as "describe and compress what happens."

What maximal description does not give us—at least not by deduction alone:

- a unique answer to “why this law?” if multiple compression-equivalent stories exist,
- a probability measure over possible laws unless the record contains evidence of a generating meta-process,
- a derivation of value, meaning, or moral obligation from physical description,
- or a final metaphysical closure that ends the recursion of “why.”

So the residual “why” remains. But it remains in a clarified form. Some “why” questions collapse into boundary conditions, universality, and selection mechanisms that *are* deducible from the record. Others persist as metaphysical choices: interpretations about ultimate ground, purpose, and value.

The proper ending, then, is not “science fails,” but “science has a boundary.” Even with perfect data, unlimited computation, and the most disciplined compression program imaginable, there remains a space in which metaphysics, theology, and philosophy continue—not as competitors to science, but as questions that science, by its own nature, does not settle by deduction.

That is the paper’s final thesis: maximal description is compatible with maximal humility.

15. Conclusion: The Inference Frontier with Complete Quantum History

The thought experiment began with an impossible gift: inert tags on every particle, perfect tracking of the universal state for all time, and unlimited analytic power. The goal was not to indulge omniscience fantasies, but to locate a boundary: to determine what such a dataset would *force* us to conclude, what it would merely *suggest*, and what it would leave untouched. Along the way, a unifying theme emerged. Once the full record is given, “law” is no longer primarily a metaphysical pronouncement. It is an inferential object: the most stable, scale-respecting compression of the history, together with a ladder of effective compressions that remain predictive under coarse-graining and perturbation.

This conclusion section draws the map one last time. Some things become rigidly inferable—dynamical generators, symmetries, causal constraints, the architecture of classical records, the arrow of time, and the renormalization ladder. Other things remain plural—ontology, subsystem carving, dual descriptions, and ultimate grounds. The right posture is therefore neither triumphalist reductionism nor skeptical despair, but disciplined humility: use the dataset to extract invariants and stable compressions, report equivalence classes honestly, and refuse to treat compression as an oracle for metaphysical closure.

15.1 What becomes rigidly inferable: dynamics, symmetries, causality, emergence

With complete quantum history in hand, many debates that feel open-ended in ordinary philosophy of physics become algorithmic—at least within the limits identified in Sections 12–13. The dataset can force conclusions in four major domains.

Dynamics (the generator, or an equivalence class of generators)
From the time-indexed state record, we can reconstruct the propagators U(t2, t1) (or quantum channels) that map the state forward, and from these infer a generator structure: Hamiltonian-like evolution if global unitarity holds, or a broader open-system generator if not. Even where identifiability is imperfect, the dataset can sharply constrain the space of admissible generators and can certify dominance relations among competing models. This makes "what is the dynamical law?" far more rigid than in our actual epistemic condition.

Symmetries and conservation structure
Once a generator (or effective generators along the renormalization ladder) is inferred, the associated symmetry and conservation patterns become recoverable as invariants: Noether-like structures, broken symmetries, emergent conservation laws, and hidden sectors detectable through invariant correlation constraints. In the full record, symmetry is not a hypothesis; it is a compressive regularity visible in how the dynamics and correlators factor.

Causality, locality, and influence constraints
Even though quantum correlations do not automatically imply causal influence, the complete record supports reconstruction of an influence graph: which degrees of freedom can affect which others, with what effective propagation limits. If the universe exhibits finite-speed influence constraints, light-cone-like structure becomes an operationally deduced regularity. If it does not, the dataset would reveal genuine nonlocal dynamics distinctly from mere entanglement. Causality becomes not a metaphysical axiom, but a feature of the inferred dynamical map and its locality structure.

Emergence as record-structure and multi-scale stability
The classical world becomes precisely characterizable as stable, redundantly encoded records in subsystems. Decoherence becomes exact entanglement bookkeeping. Objectivity becomes redundancy across environment fragments. Thermodynamic irreversibility becomes the monotone growth of inaccessible correlations under privileged coarse-grainings. And the renormalization ladder becomes an explicit derivation engine: effective laws at each scale with certified domains of validity and explicit breakdown conditions. In short: emergence becomes auditable.

Taken together, these domains mark a large region of the inference frontier that becomes rigid. Given complete history, many foundational questions become questions about which invariants

and which compressions are provably stable, rather than questions about what we can measure with limited instruments.

15.2 What remains plural: ontology, factorization, ultimate grounds

Even at this limit, the dataset does not force a single "story of what exists." Several kinds of plurality persist, and the paper treated them not as failures but as logical consequences of equivalence and computability limits.

Ontology remains plural under duality and compression-equivalence
Two descriptions can be empirically equivalent: different primitives, same invariants, same predictions, same compressive adequacy. In that situation, the dataset cannot pick one ontology without importing extra criteria that are not fixed by the record. The correct scientific output is therefore an equivalence class of descriptions plus the invariants they share.

Subsystem factorization can remain non-unique
Much of physics depends on carving the world into "systems," "environments," and "observers." But factorization of the universal Hilbert space is not uniquely given. The dataset can privilege a class of decompositions by demanding locality, robustness of record structure, and compressive adequacy—but more than one such class may remain viable. When it does, statements about "what the parts really are" remain representational choices within a constrained family.

Ultimate grounds remain outside deductive closure
Even if we infer an exceptionally compact law, we can still ask why that law and why that boundary condition. Sometimes the record may support a deeper selection mechanism (meta-law) that compresses parameters and reduces arbitrariness. Sometimes it will not. Either way, there is no guarantee that explanation terminates. Maximal description can push the "why" boundary outward, but it does not guarantee that it disappears.

In other words: the dataset can fix an enormous amount of structure, but it does not automatically convert metaphysics into deduction. It constrains metaphysics—sometimes brutally—but it does not abolish it.

15.3 The main takeaway: the universe as data, laws as compression, humility as method

The main takeaway can be stated in one sentence:

- Treat the universe as data, treat laws as the most stable compressions of that data across scales, and treat humility as the method that prevents compression from becoming an oracle.

This sentence carries three commitments.

Universe as data
The thought experiment strips away our ordinary epistemic limitations. It asks what follows if the "given" is not sparse measurements but the full record. This forces a reversal: instead of positing laws and fitting data, we let the data dictate which regularities are real, which are emergent, and which are artifacts of a chosen description.

Laws as compression
A law is not merely a formula; it is a program that summarizes the record with maximal reuse and minimal ad hoc content. That viewpoint unifies symmetry, locality, effective theory, and emergence. It also naturally produces a pluralist outcome when multiple compressions tie: laws become equivalence classes of minimal stable descriptions, not necessarily a single canonical story.

Humility as method
Humility here is not rhetorical modesty. It is a technical discipline:

- report invariants first,
- state domains of validity and breakdown conditions,
- distinguish dominance from absolute minimality,
- track underdetermination explicitly,
- and refuse to treat the best compression as metaphysical proof.

With that discipline, the compression program becomes a rigorous epistemology rather than a new idolatry of "shortest code."

15.4 Future work: formal pipelines for generator discovery and emergence auditing

If the thought experiment is to be more than philosophy, it should suggest concrete research programs—even if we can only approximate them in our real, finite-data world. The future work is therefore methodological: building formal pipelines that mirror the ideal inference tasks described throughout the paper.

Pipeline A: Generator discovery under constraints
A practical pipeline would take as input a large observational record (never complete, but increasingly rich), and output a constrained family of candidate generators. Key modules include:

- reconstruction of propagators/channels from time-series data,
- enforcement of composition consistency (semigroup/group structure),
- locality and symmetry regularization (to bias toward compressive laws),
- identifiability diagnostics (to detect underdetermination regions),
- and cross-scale consistency checks (does the inferred generator support the observed EFT ladder?).

The ideal outcome is not a single "true Hamiltonian," but a certified set of models that dominate alternatives within specified hypothesis classes.

Pipeline B: Emergence auditing
Emergence becomes auditable when we turn the qualitative claims into measurable criteria:

- pointer-structure inference: which variables become stable records?
- redundancy profiles: how many environment fragments encode the same fact?
- classicality boundaries: where does classical compression fail?
- entropy ledgers: align entanglement growth with thermodynamic irreversibility under privileged coarse-grainings.
- renormalization ladders with certified error: explicit scale-by-scale model extraction.

In practice, an emergence auditor would tell us not just that "classicality emerges," but precisely where, how, and with what stability margins.

Pipeline C: Counterfactual discipline modules
To avoid sloppy "would have happened" claims, a formal counterfactual module would:

- define permissible perturbations (physically realizable operations in the model class),
- compute sensitivity spectra and leverage maps (control knobs and bottlenecks),
- and explicitly report dependence on model-class assumptions and representation choices.

Pipeline D: Invariant-first reporting standards
Finally, the paper suggests a cultural standard: publish invariants and equivalence classes, not just favored ontologies. This is the practical response to duality and factorization ambiguity: report what the data forces and label what remains a narrative choice.

The inference frontier, then, is clear. Complete quantum history would make vast swaths of physics rigidly inferable: dynamics, symmetries, causal constraints, the structure of emergence, and the multi-scale ladder of effective laws. But it would also leave a principled remainder: plural ontologies under dualities, non-unique subsystem carvings, and the residual “why” that points beyond deduction. The best final posture is not resignation and not triumph, but a disciplined practice: let compression guide discovery, let invariants anchor claims, and let humility prevent the most powerful dataset imaginable from being mistaken for metaphysical closure.

16. References

16.1 Quantum channels and Lindblad theory

Alicki, R., and Lendi, K. (2007). Quantum Dynamical Semigroups and Applications. Springer.

Breuer, H.-P., and Petruccione, F. (2002). The Theory of Open Quantum Systems. Oxford University Press.

Choi, M.-D. (1975). Completely positive linear maps on complex matrices. Linear Algebra and its Applications, 10, 285–290.

Davies, E. B. (1976). Quantum Theory of Open Systems. Academic Press.

Gorini, V., Kossakowski, A., and Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N-level systems. Journal of Mathematical Physics, 17, 821–825.

Holevo, A. S. (2012). Quantum Systems, Channels, Information: A Mathematical Introduction. De Gruyter.

Jamiołkowski, A. (1972). Linear transformations which preserve trace and positive semidefiniteness of operators. Reports on Mathematical Physics, 3, 275–278.

Kraus, K. (1971). General state changes in quantum theory. Annals of Physics, 64, 311–335.

Kraus, K. (1983). States, Effects, and Operations: Fundamental Notions of Quantum Theory. Springer.

Lindblad, G. (1976). On the generators of quantum dynamical semigroups. Communications in Mathematical Physics, 48, 119–130.

Nielsen, M. A., and Chuang, I. L. (2000). Quantum Computation and Quantum Information. Cambridge University Press.

Rivas, A., Huelga, S. F., and Plenio, M. B. (2014). Quantum non-Markovianity: characterization, quantification and detection. Reports on Progress in Physics, 77, 094001.

Watrous, J. (2018). The Theory of Quantum Information. Cambridge University Press.

Wolf, M. M. (2012). Quantum channels and operations: guided tour (lecture notes/monograph). (Often circulated as a standard reference; cite version used.)

16.2 Hamiltonian learning and quantum system identification

Burgarth, D., Giovannetti, V., and Yuasa, K. (2009). Indirect quantum Zeno effect: Quantum control by repeated measurements. Physical Review Letters, 102, 130501. (Representative of indirect-control ideas; cite as relevant.)

Burgarth, D., Maruyama, K., Nori, F., and Osborne, T. J. (2014). Scalable quantum computation via local control of only two qubits. Physical Review A, 79, 020305. (Representative of local-control and identifiability constraints; cite as relevant.)

Chuang, I. L., and Nielsen, M. A. (1997). Prescription for experimental determination of the dynamics of a quantum black box. Journal of Modern Optics, 44, 2455–2467. (Quantum process tomography.)

Flammia, S. T., and Liu, Y.-K. (2011). Direct fidelity estimation from few Pauli measurements. Physical Review Letters, 106, 230501.

Granade, C., Ferrie, C., Wiebe, N., and Cory, D. G. (2012). Robust online Hamiltonian learning. New Journal of Physics, 14, 103013.

Haah, J., Harrow, A. W., Hastings, M. B., and Ozols, M. (2017). Sample-optimal tomography of quantum states. IEEE Transactions on Information Theory, 63, 5628–5641. (Tomography limits relevant to identification.)

Kimmel, S., Low, G. H., and Yoder, T. J. (2014). Robust calibration of a universal single-qubit gate set via robust phase estimation. Physical Review A, 92, 062315. (Gate-set calibration and robust estimation.)

Ljung, L. (1999). System Identification: Theory for the User (2nd ed.). Prentice Hall. (Classical system identification toolkit.)

Poyatos, J. F., Cirac, J. I., and Zoller, P. (1997). Complete characterization of a quantum process: the two-bit quantum gate. Physical Review Letters, 78, 390–393.

Wiebe, N., Granade, C., Ferrie, C., and Cory, D. G. (2014). Hamiltonian learning and certification using quantum resources. Physical Review Letters, 112, 190501.

16.3 Decoherence, pointer states, and quantum Darwinism

Joos, E., Zeh, H. D., Kiefer, C., Giulini, D., Kupsch, J., and Stamatescu, I.-O. (2003). Decoherence and the Appearance of a Classical World in Quantum Theory (2nd ed.). Springer.

Ollivier, H., Poulin, D., and Zurek, W. H. (2004). Environment as a witness: Selective proliferation of information and emergence of objectivity in a quantum universe. Physical Review A, 72, 042113.

Schlosshauer, M. (2007). Decoherence and the Quantum-to-Classical Transition. Springer.

Zurek, W. H. (1981). Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse? Physical Review D, 24, 1516–1525.

Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. Reviews of Modern Physics, 75, 715–775.

Zurek, W. H. (2009). Quantum Darwinism. Nature Physics, 5, 181–188.

Zwolak, M., Quan, H. T., and Zurek, W. H. (2009). Redundant imprinting of information in nonideal environments: Objective reality via quantum Darwinism. Physical Review Letters, 103, 110402.

Blume-Kohout, R., and Zurek, W. H. (2006). Quantum Darwinism: Entanglement, branches, and the emergence of classicality. (Commonly cited as foundational to Darwinism-era developments; cite journal version used.)

Riedel, C. J., and Zurek, W. H. (2011). Quantum Darwinism in a hazy environment. Physical Review A, 83, 052318. (Representative of redundancy under realistic conditions.)

16.4 Causal inference in quantum settings and indefinite causal structure

Pearl, J. (2000). Causality: Models, Reasoning, and Inference. Cambridge University Press.

Spirtes, P., Glymour, C., and Scheines, R. (2000). Causation, Prediction, and Search (2nd ed.). MIT Press.

Peters, J., Janzing, D., and Schölkopf, B. (2017). Elements of Causal Inference: Foundations and Learning Algorithms. MIT Press.

Leifer, M. S., and Spekkens, R. W. (2013). Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. Physical Review A, 88, 052130.

Oreshkov, O., Costa, F., and Brukner, Č. (2012). Quantum correlations with no causal order. Nature Communications, 3, 1092.

Chiribella, G., D'Ariano, G. M., and Perinotti, P. (2009). Theoretical framework for quantum networks. Physical Review A, 80, 022339. (Quantum combs/processes.)

Henson, J., Lal, R., and Pusey, M. F. (2014). Theory-independent limits on correlations from indefinite causal order. New Journal of Physics, 16, 113043.

Costa, F., and Shrapnel, S. (2016). Quantum causal modelling. New Journal of Physics, 18, 063032.

Fritz, T. (2012). Beyond Bell's theorem: correlation scenarios. New Journal of Physics, 14, 103001. (Generalized causal/correlation frameworks.)

16.5 Emergent spacetime/geometry from entanglement and holography

Maldacena, J. (1998). The large-N limit of superconformal field theories and supergravity. Advances in Theoretical and Mathematical Physics, 2, 231–252. (Often cited as 1997 preprint; cite journal year/version used.)

Ryu, S., and Takayanagi, T. (2006). Holographic derivation of entanglement entropy from AdS/CFT. Physical Review Letters, 96, 181602.

Hubeny, V. E., Rangamani, M., and Takayanagi, T. (2007). A covariant holographic entanglement entropy proposal. Journal of High Energy Physics, 07, 062.

Van Raamsdonk, M. (2010). Building up spacetime with quantum entanglement. General Relativity and Gravitation, 42, 2323–2329. (Also widely circulated as an essay/preprint; cite version used.)

Swingle, B. (2012). Entanglement renormalization and holography. Physical Review D, 86, 065007.

Pastawski, F., Yoshida, B., Harlow, D., and Preskill, J. (2015). Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence. Journal of High Energy Physics, 06, 149.

Almheiri, A., Dong, X., and Harlow, D. (2015). Bulk locality and quantum error correction in AdS/CFT. Journal of High Energy Physics, 04, 163.

Jafferis, D. L., Lewkowycz, A., Maldacena, J., and Suh, S. J. (2016). Relative entropy equals bulk relative entropy. Journal of High Energy Physics, 06, 004.

Jacobson, T. (1995). Thermodynamics of spacetime: The Einstein equation of state. Physical Review Letters, 75, 1260–1263.

Susskind, L., and Maldacena, J. (2013). Cool horizons for entangled black holes. Fortschritte der Physik, 61, 781–811. (ER=EPR theme.)

16.6 Renormalization group and effective field theory

Kadanoff, L. P. (1966). Scaling laws for Ising models near Tc. Physics, 2, 263–272.

Gell-Mann, M., and Low, F. E. (1954). Quantum electrodynamics at small distances. Physical Review, 95, 1300–1312.

Wilson, K. G. (1971). Renormalization group and critical phenomena. I. Renormalization group and the Kadanoff scaling picture. Physical Review B, 4, 3174–3183.

Wilson, K. G., and Kogut, J. (1974). The renormalization group and the epsilon expansion. Physics Reports, 12, 75–199.

Polchinski, J. (1984). Renormalization and effective Lagrangians. Nuclear Physics B, 231, 269–295.

Weinberg, S. (1979). Phenomenological Lagrangians. Physica A, 96, 327–340.

Georgi, H. (1993). Effective field theory. Annual Review of Nuclear and Particle Science, 43, 209–252.

Kaplan, D. B. (2005). Five lectures on effective field theory. (Widely cited pedagogical notes; cite version used.)

Peskin, M. E., and Schroeder, D. V. (1995). An Introduction to Quantum Field Theory. Westview Press.

Shankar, R. (1994). Renormalization-group approach to interacting fermions. Reviews of Modern Physics, 66, 129–192.

Cardy, J. (1996). Scaling and Renormalization in Statistical Physics. Cambridge University Press.

Goldenfeld, N. (1992). Lectures on Phase Transitions and the Renormalization Group. Addison-Wesley.

16.7 Algorithmic information theory, MDL, and computability limits

Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society, 42, 230–265.

Church, A. (1936). An unsolvable problem of elementary number theory. American Journal of Mathematics, 58, 345–363.

Rice, H. G. (1953). Classes of recursively enumerable sets and their decision problems. Transactions of the American Mathematical Society, 74, 358–366.

Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. Problems of Information Transmission, 1, 1–7.

Solomonoff, R. J. (1964). A formal theory of inductive inference. Information and Control, 7, 1–22 and 224–254.

Chaitin, G. J. (1975). A theory of program size formally identical to information theory. Journal of the ACM, 22, 329–340.

Levin, L. A. (1974). Laws of information conservation (also cited via universal search / Levin complexity). (Cite standard archival translation/source used.)

Rissanen, J. (1978). Modeling by shortest data description. Automatica, 14, 465–471. (Foundational MDL.)

Li, M., and Vitányi, P. (2008). An Introduction to Kolmogorov Complexity and Its Applications (3rd ed.). Springer.

Cover, T. M., and Thomas, J. A. (2006). Elements of Information Theory (2nd ed.). Wiley.

Martin-Löf, P. (1966). The definition of random sequences. Information and Control, 9, 602–619.

Bennett, C. H. (1988). Logical depth and physical complexity. In The Universal Turing Machine: A Half-Century Survey (ed. Herken). Oxford University Press.

The author has published over 4,800 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.

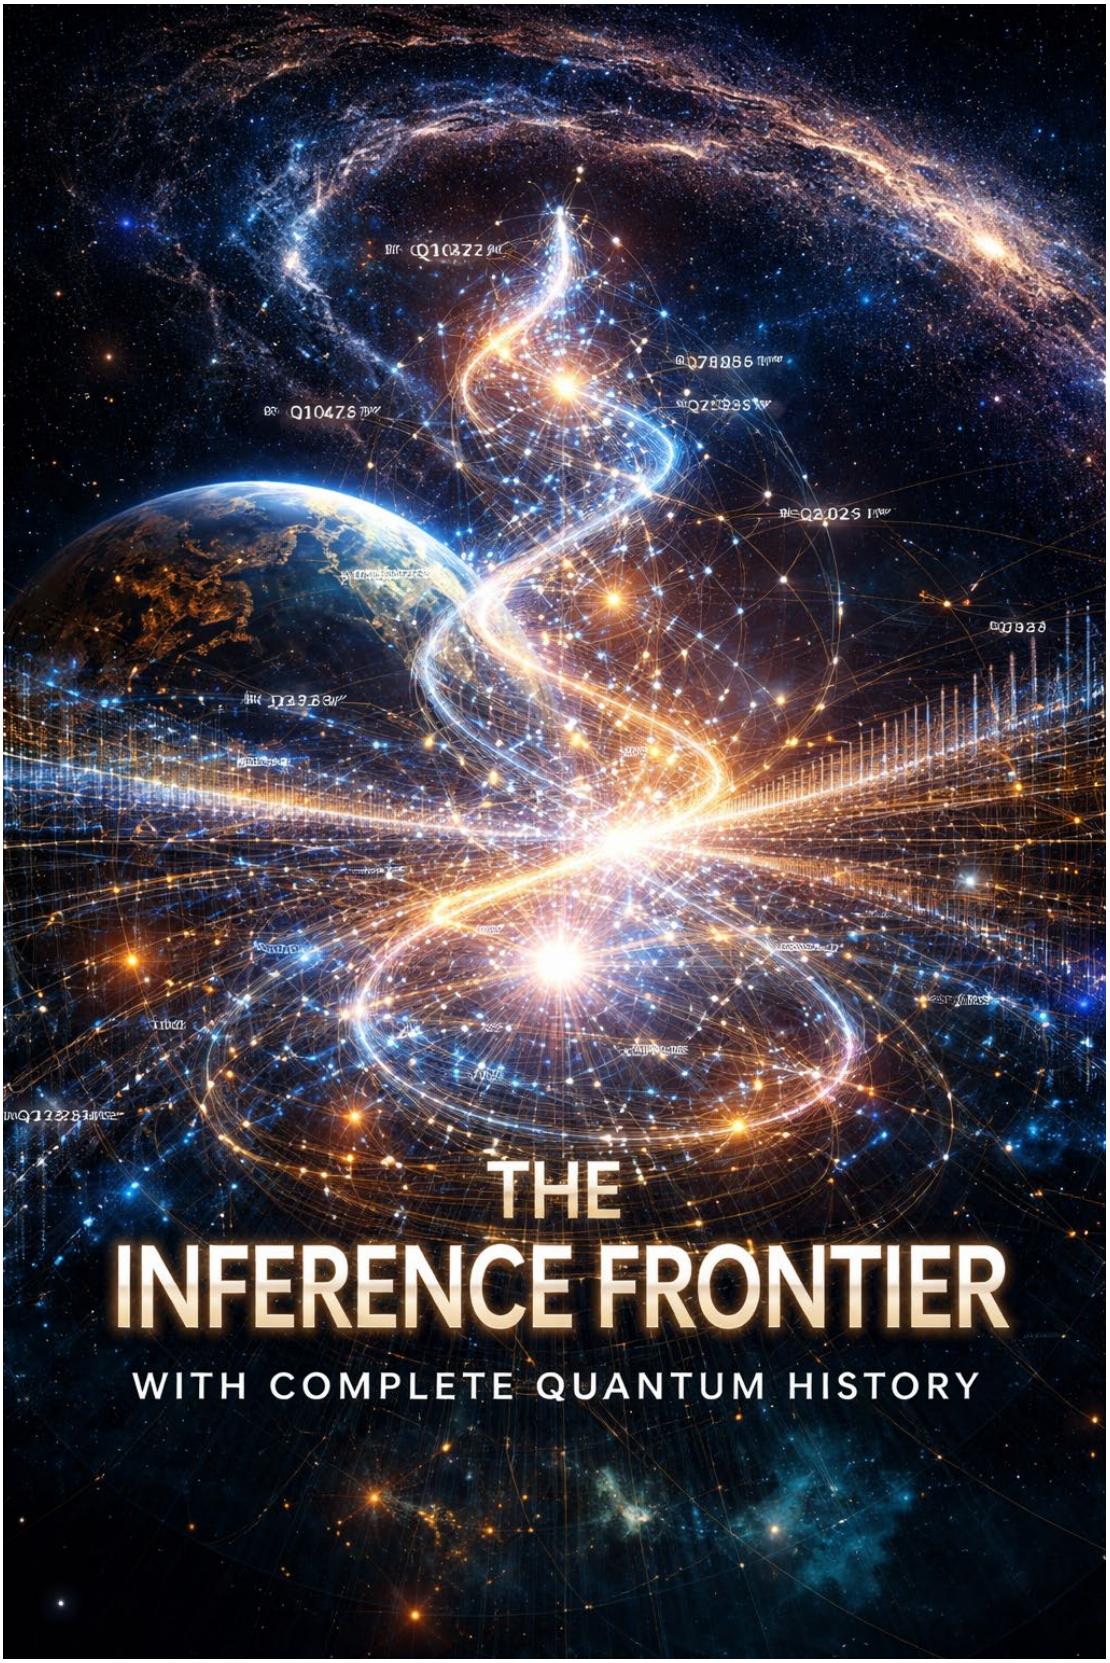
THE
INFERENCE FRONTIER
WITH COMPLETE QUANTUM HISTORY